Techdebate **NET NEUTRALITY -** Is it a necessity? Two experts weigh in. **PAGE 20**

NETWORKWORLD

THE CONNECTED ENTERPRISE ≡ MARCH 21, 2011

VOLUME 28 NUMBER 6 $5.00

CLEAR CHOICE TEST
CISCO ANYCONNECT SECURE MOBILITY

Cisco sets the bar for mobile security

Cisco integrates always-on client, VPN/firewall and Web security gateway. **Page 23**

Cloud-focused HP not backing down from IBM, Cisco

BY JOHN GALLANT, IDGE, AND ERIC KNORR, INFOWORLD

LEO APOTHEKER Q&A

A DAY after HP CEO Leo Apotheker outlined his strategic vision for HP — a plan chockfull of new cloud offerings (see page 13) — he sat down with us to share his thoughts on everything from why he thinks HP is better positioned than IBM to help customers deliver on the promise of cloud to how HP is tackling mobility to what it's like competing with Cisco.

How will you help companies get to a hybrid public/private cloud model?

There will be as many combinations between traditional and on-premise private clouds, public clouds, semipublic clouds, as there are enterprises. One of the reasons things

▶ *See* **HP,** *page 12*

On the company dime: Rogue game server admins tell all

BY PAUL MCNAMARA

BACK IN January, Scandinavian gamers hijacked a New Hampshire medical center's server to host "Call of Duty: Black Ops" sessions. When asked about that incident, Stephen Heaslip of the gamer site *Blue's News* told *Network World* that hackers are not the most likely individuals to commandeer corporate servers for illicit gaming: Such appropriations are more often the work of IT administrators. When asked if he could put us in touch with some of these rogue game server admins, Heaslip posted a call to his readership — and four volunteers stepped forward.

▶ *See* **Rogue,** *page 14*


BladeCenter E
IBM

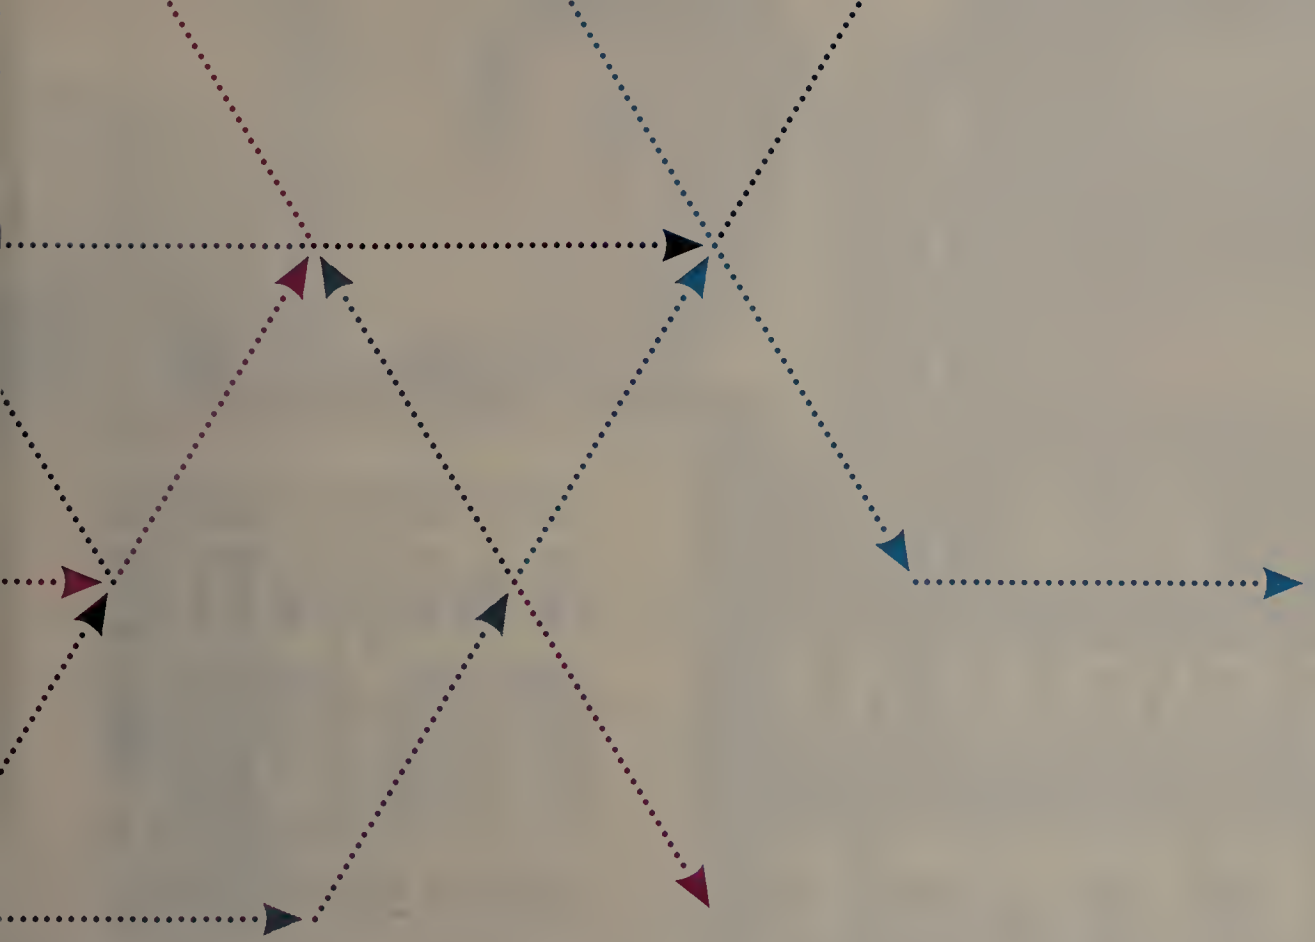

The difference between networking and not working.

Some systems require you to reconfigure your network infrastructure to match their standards. Not the IBM BladeCenter® with Intel® Xeon® processors. It offers a broad range of networking technologies—including some of the most advanced virtualization solutions in the industry. So you can choose the one that works best with your infrastructure. And IBM BladeCenter can save you up to 40% on networking costs versus competitive offerings.[1]

Take 10 minutes to see for yourself.

Learn how you could achieve a 3-month ROI on your migration with our Systems Consolidation Tool. Visit ibm.com/systems/blade

1. The 40% cost savings are based on a comparison of the acquisition costs of 10 current generation HP rack optimized solutions (i.e. DL380 G7 Proliant with 10 GbE Ethernet and Fibre Channel infrastructure) to 10 current generation IBM BladeCenter and HS22 systems with converged fabric solutions from Brocade. See www-03.ibm.com/systems/bladecenter/hardware/openfabric/fcoe.html. The IBM solution includes chassis infrastructure. Pricing utilizes publicly available pricing per port for ToR ethernet and FC switching infrastructure as of Jan 2011. The 40% networking hardware costs savings result from eliminating separate Ethernet and Fibre Channel cards and switches in the deployment of an IBM BladeCenter FCoE solution for 10 servers and associated networking hardware in comparison to the HP solution. IBM, the IBM logo, ibm.com and BladeCenter are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml. Intel, the Intel logo, Xeon and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

1&1® WEB HOSTING

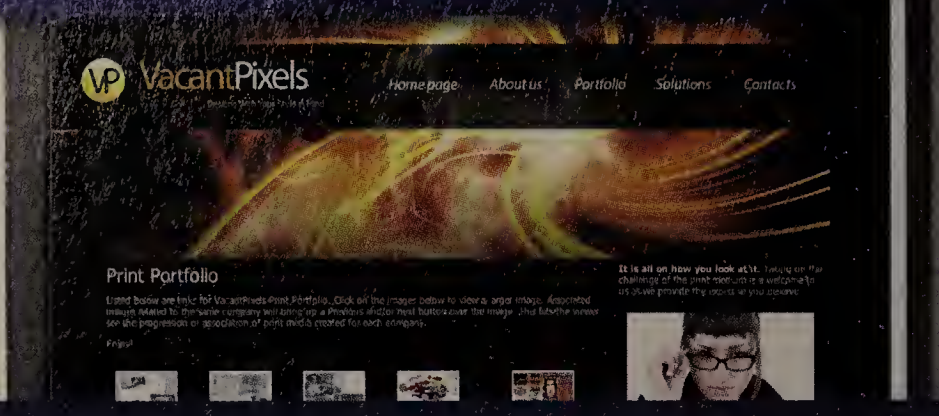

PROFESSIONAL WEBSITES

As the world's largest web host, we know the developer features you need in a hosting package!

.com
.info .org
.net

Domains Included
All hosting packages include domains, free for the life of your package.

Unlimited Traffic
Unlimited traffic to all websites in your 1&1 hosting package.

Developer Features
Extensive language support with PHP 5/6 (beta) with Zend Framework and git version management software.

Online Marketing Tools
SEO tools to optimize your website.
1&1 Webstatistics makes it easy to monitor your progress.

Green Data Centers
We're committed to hosting your site with a minimal impact on the environment.

1&1® HOSTING PACKAGES

6 MONTHS FREE!*

OFFER EXTENDED!

1&1® BUSINESS PACKAGE:

- 3 Included Domains
- Private Domain Registration
- 250 GB Web Space
- UNLIMITED Traffic
- NEW: Version Management Software (git)
- 2,500 E-mail Accounts
- 50 MySQL Database (100 MB)
- 25 FTP Accounts
- E-mail Marketing Tool
- 24/7 Toll-free Customer Support

~~$0.99~~ per month*

Need more domains?
.info domain only $0.99 first year*
.com domain only $4.99 first year*

More special offers available on our website!

1&1®

MEMBER OF united internet

Get started today, call 1-877-GO-1AND1 **www.1and1.com**

*Offers valid for a limited time only. 12 month minimum contract term applies for web hosting offers. Setup fee and other terms and conditions may apply. Domain offers valid first year only. After first year, standard pricing applies. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

FROM THE EDITOR | JOHN DIX

The change imperative

Some 400 of your peers gathered at *Network World*'s IT Roadmap show in Chicago last week to discuss everything from cloud plans to iPad support, flat network options and Windows 7 migration.

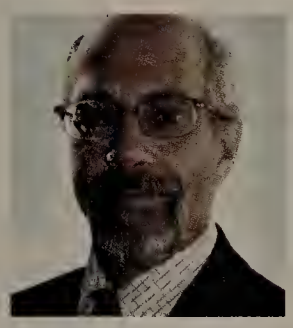

Key take-away: Things are changing fast. Consider cloud. In a panel discussion with Chad Eckes, CIO of Cancer Treatment Centers of America, Karthikeyan Chakkarapani, IT director of technology solutions and operations at American Hospital Association, Rob Zelinka, former director of infrastructure at TTX and Tomasz Chowanski, IT leader of shared-services security architecture at GE Capital; I asked if anyone had solid cloud plans. As it turns out, all of them are already doing some cloud computing.

So I asked the audience to show by raising their hands how many would be doing something with cloud computing this year. Three-quarters of them waved.

How about allowing employees to access social media sites? All of the panel members said "check," and three-quarters of the audience agreed. When I asked the audience if their companies allowed it a year ago, half the hands dropped.

Anyone still trying to fight off iPads and other employee-owned gadgets? The panel said no, but added qualifiers. Chakkarapani said Citrix and other tools make it pretty easy to accommodate tablets, but the speakers were in general agreement about needing controls that let you manage the devices, guard against data loss, and wipe them if you need too.

Chowanski said that once you spell all of that out for users, many think twice about seeking net access. People don't like the idea of the company being able to see everything on their gizmos and, worse, the prospect of losing their photos, videos, music and other files.

Speaking of network endpoints, the session on migrating to Windows 7 and Office 2010 drew a standing room-only crowd. With XP set to sunset in 2014, W7 migration will really heat up this year because implementation can take 14 to 18 months in a big shop, said speaker Sevan Muradian, senior product marketing manager at Dell KACE.

A detailed inventory is critical before you get going, Muradian said, and audience polling proved what he attested: Most people don't have a good handle on the hardware in their shops, let alone the desktop software floating around.

One big pitfall to be aware of: the need to do regression testing on all the Excel macros that departments have built to support their processes.

As young as the year is, it looks to be shaping up as one for great change in IT. Join the conversation at an IT Roadmap coming to a town near you (Denver next month and Boston in June).

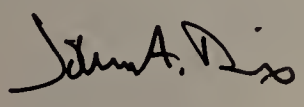

COVER ILLUSTRATION: STEPHEN SAUER

inside

MARCH 21, 2011

peersay

Stallman is clear about what he stands for

WHETHER YOU AGREE with Richard Stallman or not, you know where he stands when it comes to software freedom. And without the Free Software Foundation's GNU tools, Linux might not have succeeded as quickly as it did, so I can understand his interest in having GNU properly credited (re: "Cell phones are 'Stalin's dream,' says free software movement founder"; tinyurl.com/6xb724e).

As for Hurd, even though it is still an active project, you have to wonder why it was never completed in a more timely manner. Linus Torvalds managed to hack out the Linux kernel in less than 18 months and he invited others to join in the effort. Thus was born the open source software meritocracy where individuals with software skills could contribute to interesting projects and be recognized by the open source community for doing good work. The open source software development model is now widely accepted as a legitimate alternative to proprietary software. I'd say we are better off on the whole with the opportunity to run both free and open source software.

Richard Stallman deserves our respect for being the voice for software freedom. He talks the talk and walks the walk.

twessels

Tablets and smartphones in enterprise

THERE ARE TWO types of risk. One, to the organization, of sensitive content being exposed if the device is lost, hacked, or otherwise compromised. In some cases there are financial penalties for this, as well as costly notification practices that need to be complied with if it involves any customer data (re: "Wells Fargo says no to personal smartphones and tablets, period"; tinyurl.com/4aw4af9).

The other is to the employee. In the event of a legal action involving anything they may have been involved in, or a data call to "... produce any/all records related to XYZ," the employee's device may be subject to search. This could risk exposing their personal data, including passwords, contacts, browser history and other things they may not want their employer or others to have access to.

Commingling business/personal content and activity just plain isn't good sense. Even a one-person consulting business keeps its personal and business financial assets/accounts independent of each other; why doesn't it make the same sense to keep your information assets independent?

Larry

Anonymity, privacy, control and money

WHERE YOU STAND on this issue depends primarily on where you stand economically and to some extent morally. Those who are on top or who control or have accepted being controlled and dealt with like a tool as a way of life will undoubtedly vote for the repeal of anonymity. Those who are not on the top, the abused, the downtrodden, the revolutionary, will take the opposite route (re: "4chan founder moot: 'Anonymity is authenticity,' Zuckerberg 'wrong'"; tinyurl.com/5tyfm7t).

> Commingling business/personal content and activity just plain isn't good sense.

If anonymity was not allowed on the Internet, every move you made could be captured and sold for a profit. There would be no privacy, since once they knew you and you couldn't change your identity you'd become an open, indefensible target for political and economic control.

On the other hand, too much privacy is harmful too. Since anonymity and privacy allow the full individual's psyche to emerge, there must be control, or you'll get sexual predators, bullies, etc. to come out.

The answer lies in between. How we achieve it will determine if the Internet is just a utility system for exploiting the lower classes of humanity or a conduit for the growth of mankind. So far, unfortunately it looks like its just going to be a utility system for the powerful to exploit the weak.

Richard the Mongoose

NETWORKWORLD

492 Old Connecticut Path, Framingham, MA 01701-9002
Main Phone: (508) 766-5301
E-mail: firstinitial_lastname@nww.com
Editorial Calendar: http://tinyurl.com/39sf649

EDITORIAL
Editor in Chief: John Dix
Online Executive Editor, News: Bob Brown
Executive Online Editor: Jeff Caruso
Executive Features Editor: Neal Weinberg
Community Editor: Julie Bort
Multimedia Programming Director: Keith Shaw

NEWS EDITORS
Online News Editor: Michael Cooney
Online News Editor: Paul McNamara
Online Associate News Editor: Ann Bednarz

REPORTERS
Jon Brodkin, Senior Editor
John Cox, Senior Editor
Jim Duffy, Managing Editor
Tim Greene, Senior Editor
Carolyn Duffy Marsan, National Correspondent
Ellen Messmer, Senior Editor
Brad Reed, Senior Writer

PRINT LAYOUT/WEB PRODUCTION
Managing Editor: Ryan Francis
Web Producer: Melissa Andersen

DESIGN
Executive Art Director: Mary Lester
Associate Art Director: Stephen Sauer

NETWORK WORLD LAB ALLIANCE
Joel Snyder, Opus One; John Bass, Centennial Networking Labs; Barry Nance, independent consultant; Thomas Henderson, ExtremeLabs; David Newman, Network Test; James Gaskin, Gaskin Computing Services; Craig Mathias, Far Point Group

OFFICE MANAGEMENT
Editorial Operations Manager: Cheryl Crivello
Office Manager, Editorial: Pat Josefek

SUBSCRIPTIONS
Phone: (877) 701-2228
E-mail: nww@omeda.com
URL: www.subscribenww.com

REPRINTS
(717) 505-9701, ext. 129

idg Enterprise

IDG ENTERPRISE
CEO: Mike Friedenberg
Group Publisher: Bob Melk
Chief Content Officer/SVP: John Gallant

bits

Notorious spamming botnet takes a fall

FOR MORE THAN 24 hours last week, it was a question very few security experts could answer: Who had knocked the world's worst spam botnet offline? After infecting close to a million computers and spamming out as many as 30 billion unwanted e-mail messages a day, the Rustock botnet went silent last Wednesday. Now we know why: A small group of computer researchers, backed by Microsoft's lawyers and international law enforcement, executed a number of surgical strikes on the botnet. Hitting it as if it were the mythical Hydra, they cut off Rustock's heads — its command-and-control servers — and scorched them to keep them from growing back. Now Microsoft is helping to clean up infected computers before Rustock's owners have a chance to regain control of their botnet. ***tinyurl.com/4a5xtj9***

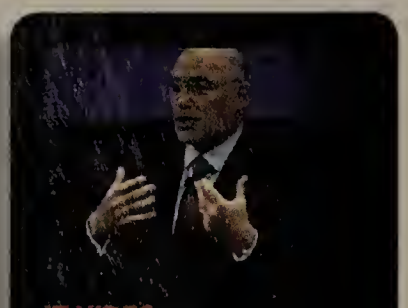

IT VIDEO

HP CEO bets company on cloud

HP also plans to open a marketplace that will include both applications and cloud-based services for enterprises, small businesses and consumers. ***tinyurl.com/69ttbg8***

Facebook likes microservers

FACEBOOK IS bucking the trend toward server virtualization and is interested in microservers for inexpensive growth and quick failover, according to lab director Gio Coglitore. The social networking giant came out in support of Intel's plans for an expanded lineup of processors for microservers, which are small, low-power, one-processor servers that can be packed into a data center more densely than rack or blade servers. At an event last week, Intel said it would introduce four new chips for microservers this year and next, ranging from a 45-watt Xeon to an Atom-based processor that consumes less than 10 watts. All will have server-class features, such as 64-bit compatibility and ECC (error-correcting code) memory. Facebook has tested microservers in production and is interested in the architecture for its massive data centers, Coglitore says, but the key is inclusion of these new server features. ***tinyurl.com/4q9g858***

Firefox 4 coming this week

MOZILLA IS expected to ship the final version of Firefox 4 on Tuesday. Originally scheduled to ship last November, Firefox 4 will wrap up a development cycle that started in February 2010 with several previews, but began in earnest last July when Mozilla released the first of what would eventually be a dozen betas. Last week, Mozilla developers called the current Release Candidate good enough to ship as the final. "Today's triage session concluded with all systems go for a Firefox 4 launch on March 22," said Damon Sicore, Mozilla's senior director of platform engineering. Firefox 4 features a new tab manager, dubbed "Panorama," supports GPU acceleration to boost page composition speeds, and boasts an overhauled interface that resembles Chrome's and IE9's minimalist designs. ***tinyurl.com/4cne6sl***

How to put the brakes on Google Apps upgrades

GOOGLE WILL start letting apps administrators delay the delivery of upgrades to their domains to give them a chance to prepare themselves and their users for interface or functionality changes. Until now, Google has transparently pushed out enhancements to its Apps suite on a rolling basis as soon as they were tested and deemed ready for prime time, just like it does for its consumer applications and sites. Apps administrators, especially those in large enterprises, said they wanted a heads-up about upcoming changes and a time buffer before the upgrades go live on their domains. Industry analyst Rebecca Wettemann calls Google's new "Scheduled Release" track a good idea. "It shows Google is working to make their apps more digestible for enterprise organizations," she said. ***tinyurl.com/4jq8xc7***

Happy 5th birthday, Twitter

TWITTER IS celebrating its 5th birthday this month with a fresh set of stats about its growth and usage, such as the fact that it took more than 3 years for users to send the first 1 billion tweets, a feat now accomplished every week. Over the past year, the average number of Twitter messages sent per day has increased from 50 million in March 2010 to 140 million this month. Currently, an average of 460,000 Twitter accounts are created every day, while the number of Twitter mobile users has spiked 182% in the past year. ***tinyurl.com/4byzb4r***

Iyad Tarazi
VICE PRESIDENT OF NETWORK DEVELOPMENT & ENGINEERING
SPRINT

Iyad Tarazi is responsible for the development and integration of new products and technologies within Sprint's wireless and wireline networks and for the engineering of the RF and core networks.

FOR MORE INFORMATION:
please visit **www.Sprintv6.net**

Sprint

NETWORKWORLD
Custom Solutions Group

Gearing Up for IPv6

CIOs start down the long road to IPv6 with a little help from Sprint.

Given the growth rate of the Internet, IPv6 adoption is a much needed reprieve. Sprint's Tarazi explores what it's going to take to make the transition.

What is the promise of IPv6?
The reason behind IPv6 is that it allows a lot more devices to connect to the Internet. When IPv4 was created it envisioned a much smaller Internet. Today, though, studies indicate as many as 1 trillion devices could be connected to the Internet by 2013. Machines will talk to machines, and we are heavily involved in driving this activity. The possibilities are limitless, but an example today includes wirelessly embedded devices for remotely monitoring chronic health conditions. There will be a need for a lot more addressing space, and IPv6 specifically solves that issue. The good news is IPv6 is already here.

What are the cold realities of transitioning to IPv6?
The biggest challenge is the Internet itself. The adoption of Internet protocol has been so pervasive that it is embedded in many of our business and operational systems, devices and applications. It will be difficult to find and work out a plan for those thousands of network touch-points within the enterprise. The analysis required will be quite similar to what enterprises experienced with Y2K—but on a bigger scale—in that the change is somewhat easy but it's embedded in so many places that execution can prove challenging. But unlike Y2K, there will not be a definitive cut-off date; enterprises that haven't transitioned will experience the impact more slowly.

What can CIOs expect to experience at the tipping point?
Thus far, there has been wide adoption by carrier and equipment providers. In the future, just about everything new that you buy will be on an IPv6 protocol. At the tipping point when the volume of IPv6 traffic is significant enough, enterprises will see shrinking support for IPv4 issues as manufacturers focus on IPv6 enabled devices, and translation solutions will become cost ineffective due to scale. So there has been a lot more talk in the industry about how to create transition equipment, such as dual-stack and translation devices. This will enable you to move forward without the significant cost structure associated with a major "rip and replace."

How can CIOs best manage migration?
First, you need to set up a program office, where you do the planning activities and set resources aside to map out your enterprise systems, interface points and partners, carriers, and suppliers. Then you need to develop a strategy for transition. You should create a gateway-type architecture to translate and protect older IPv4 systems while they are being updated to IPv6. Dual-stack capability in systems and user platforms is an important part of a gradual transition. It will take time to translate the puts and takes of your network into requirements for your partners and providers, and your analysis and planning should start right now. For most businesses full transition will need to happen within typical business development and equipment replacement cycles, or about three to five years. You won't just wake up one day and decide what to do.

How has Sprint prepared for the successful adoption of IPv6?
We've been at this for years. We have already added IPv6 to our core system and have begun the process of migrating most of our edge equipment, so we can provide more customer- and government-facing systems that can talk IPv6. As we convert, we've identified and prioritized those systems that need IPv6 first, and have deferred those where IPv6 is not relevant in the network. Finally, we are supporting customers with dual-stack mode so they can communicate with the IPv6 protocol. We've learned a lot throughout this complex process and are happy to transfer that knowledge—as adviser, helper, or educator—to anyone that needs it.

Rely on one network right where your team needs it.

Sprint Global MPLS gives you the upper hand by converging voice, video and data on a single IP-based network. Sprint Global MPLS also gives you best-in-class network performance, with industry-leading SLAs and Class of Service at no additional charge to get you started.

1-866-653-1056 sprint.com/convergence

Sprint

The Now Network™

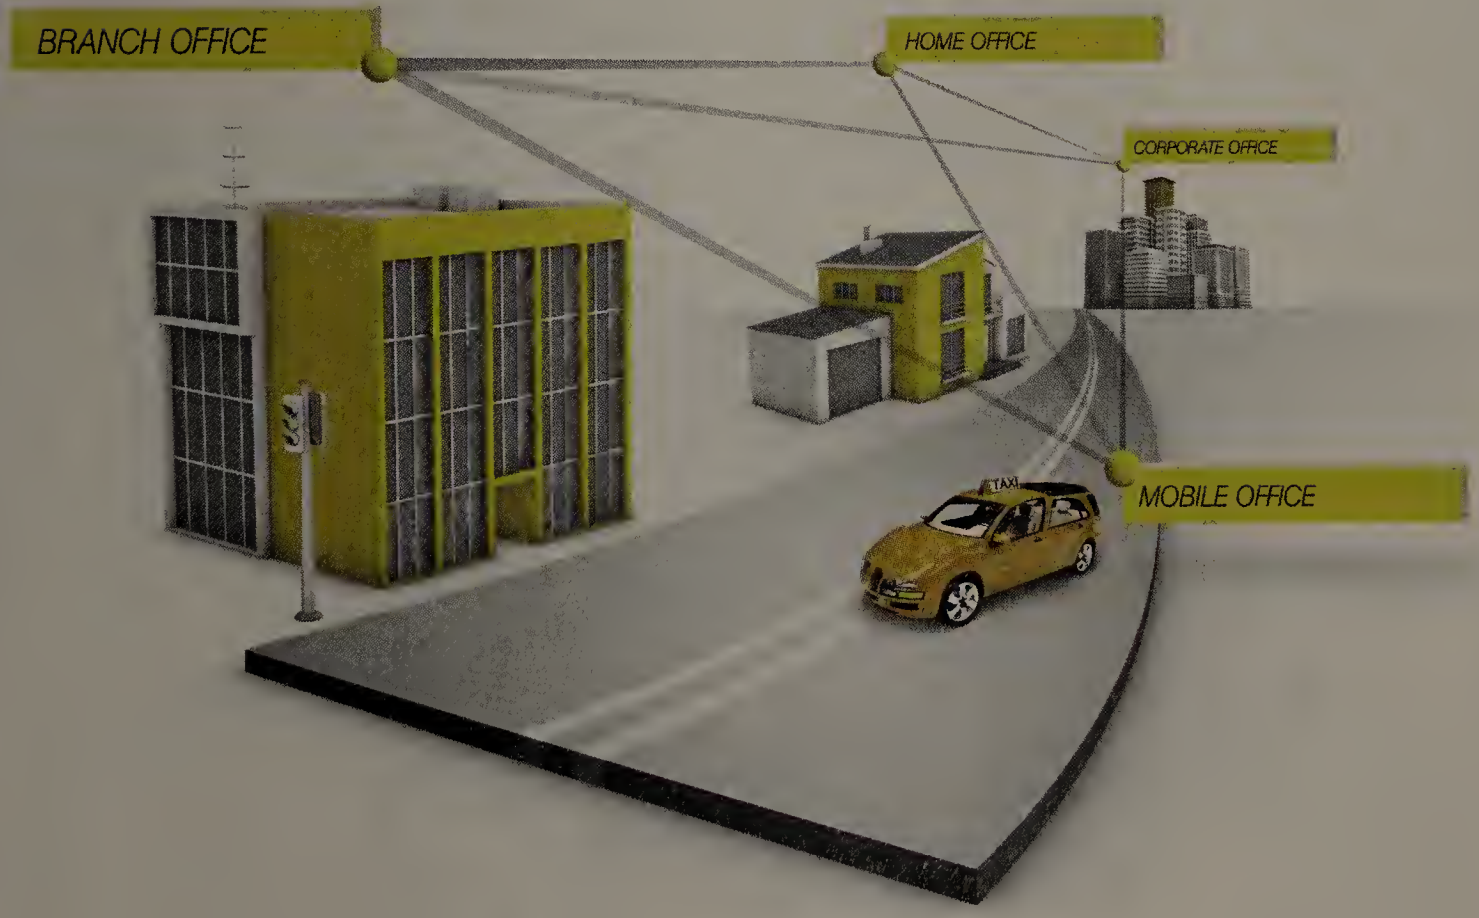

Winner of the Frost & Sullivan North American Product Leadership Award
for MPLS Service Level Agreements for Business Customers—2010

GOOD BAD UGLY

IE9 in demand

good

MICROSOFT BOASTED that Internet Explorer 9 was downloaded 2.35 million times, or 27 times per second, during its first 24 hours of availability last week. Microsoft has always loved to talk about how many copies of software it can move per second. Last year, Microsoft bragged about selling 7 copies of Windows 7 every second. The question is whether Microsoft can keep up the momentum. Internet Explorer is still the most widely used browser, but market share has slipped consistently in the past few years because of the growing popularity of Mozilla Firefox and Google Chrome.

iPad 2 tablets too popular

bad

APPLE HAS a problem so many companies would like to have: its new iPad 2 tablets are proving to be so popular that many stores are selling out and now people ordering online are being forced to wait four weeks or more for delivery. Apple watchers keep adjusting their sales projections, with analysts pegging sales during the first weekend at as high as 1 million. What's more, the disaster in Japan could put a squeeze on NAND flash memory supplies needed by Apple for its iPads and other devices.

School is open

ugly

EDUCATIONAL INSTITUTIONS and social networks are the worst when it comes to leaving their Web sites exposed to known vulnerabilities such as cross-site scripting and SQL injection, according to a study by WhiteHat Security. According to its 11th annual Web Site Security Statistics Report, 71% of schools have unpatched software vulnerabilities on their Web servers all the time, while 58% of social networking sites always have such vulnerabilities. By contrast, 14% of health care organizations and 16% of banks have unpatched vulnerabilities all the time.

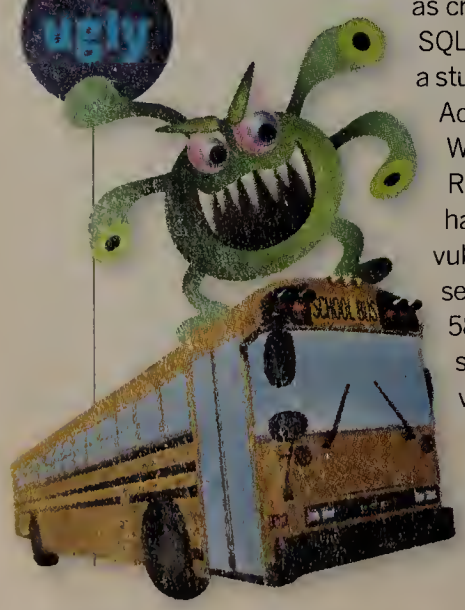

DAMIR SAGOLJ / REUTERS

Japanese quake may shorten days

JAPAN'S MARCH 11 earthquake may have shifted the Earth's mass enough to change its rotation and result in shorter days, spurring changes in computer time-keeping. The 9.0-magnitude quake could have shortened days by up to 1.8 microseconds, argues research scientist Richard Gross of NASA's Jet Propulsion Laboratory in Pasadena, Calif. A microsecond is one-millionth of a second. While humans obviously won't notice the change, it will have to be recognized in the official time-keeping systems used for reconciling computer time with solar time. ***tinyurl.com/4qddnuv***

Party lines, 2011 style

TALK

WOULDN'T IT be great if you could press a button on your smartphone and broadcast to your friends or fellow workers walkie-talkie style? This kind of push-to-talk communications is now possible using Twisted Pair Solutions' new Wave Connections service. The hosted offering lets smartphone users sign up, download a small Wave Connections client and then invite other smartphone users to do the same. When they do, all users in the group have instant, push-to-talk access (you designate any button on your phone as the PTT button) with everyone else in the group. The app is initially available for RIM BlackBerry, Microsoft Windows Mobile and Windows CE phone users, with Google Android support on the way, followed by Windows Phone 7. ***tinyurl.com/4tqnblq***

Missing: 9 server drives, lots of customer data

HEALTH NET, a provider of managed health care services, has been alerting some 1.9 million customers that nine disk drives containing personal and health data were recently discovered missing from a data center managed by IBM in Rancho Cordova, Calif. An initial probe found the drives contained names, addresses, Social Security numbers, financial information and health data of current and former Health Net members, employees and health care providers. Health Net said it will offer two years of free credit monitoring services to the affected individuals. ***tinyurl.com/4nwd3an***

DIY network in the cloud

AMAZON WEB Services (AWS) has added a number of networking features to its Virtual Private Cloud (VPC) offering, allowing users to build data centers in the cloud that can be private, accessed from the Internet or both. The added features let users build a virtual network architecture, with full control over routing and subnets. ***tinyurl.com/4hdegnl***

AMD Opteron™ 64

UNLOCK

the possibilities of change-ready computing with the power of convergence.

HP Converged Infrastructure unlocks what's next with virtualization on HP BladeSystem.

Server virtualization has cut costs, but you can do more with the simplest way to connect to networks—HP Virtual Connect.

With HP BladeSystem featuring HP ProLiant G7 servers powered by AMD Opteron™ 6100 Series processors, HP Virtual Connect lets you wire once so you can make changes in minutes rather than days, cutting complexity and saving time.

HP Virtual Connect can deliver up to:

- 95% less network sprawl at the server edge*
- 65% cost savings on equipment*
- 40% lower power and cooling costs*

Unlock the possibilities with a free copy of *HP Virtual Connect for Dummies*®.
>> hp.com/go/complexity6

HP ProLiant BL465c G7 server

- Two twelve-core AMD Opteron™ 6100 Series processors installed
- 16GB of memory installed; expandable to 256GB
- HP Smart Array P410i Controller with 1GB Flash Backed Write Cache installed
- One integrated NC551i Dual Port 10Gb FlexFabric Converged Network Adapter
- Up to two HP hot plug small form factor SAS, SATA, or Solid State drives

$4,509 (Save $777)†

Lease for just $110/mo.†

SmartBuy (PN: 630444-S01)

*For details an claim substantiations, visit hp.com/go/complexity6

†Prices shown ore HP Direct prices, reseller ond retoil prices moy vary. Prices shown are subject to change, ond do not include applicable state and local taxes or shipping to recipient's address. Offers cannot be combined with any other offer or discount ond are good while supplies last. All featured avoiloble offers in U.S. only. Savings based on HP published list price of configure-to-order equivalent (HP ProLiont BL465c G7 server: $5,286 - $777 = SmortBuy price $4,509). Financing available through Hewlett-Packard Financiol Services Campony and its subsidiaries (HPFSC) to qualified commercial customers in the U.S. and is subject lo credit approval and execution of standard HPFSC documentation. Prices shown are based on o lease 48 months in term with a fair market value purchase option ot the end of the term, ond are valid through December 31, 2011. Other charges and restrictions may apply. HPFSC reserves the right to change or cancel this program at any time without natice.

► **HP,** *from page 1*

don't go straight into the cloud is the legacy of applications. Some of these applications would be very hard to move into the cloud if you don't want to provoke a rainstorm and [have the cloud collapse].

HP has a lot of experience helping customers make these decisions, make the trade-offs, and then help people move into these hybrid environments. We create hardware and software to manage hybrid environments. Some of our technology allows people to have a complete end-to-end vision of all of these mixed architectures and operate them as one.

In the broad portfolio of capabilities that you've presented, it seems to overlap almost 100% with what IBM is doing. How do you intend to differentiate your strategy?

I would qualify it slightly differently: IBM overlaps 100% with us. HP has been doing these things for years — we didn't really call them out this way — but this is nothing really that revolutionary or new.

We have [several] strategic advantages over IBM. One is we understand the consumer business, so therefore we understand the endpoint devices. And that is a huge advantage, which IBM has given away when they sold their PC business to Lenovo. Secondly, we have deep insight into security and manageability, which helps us to secure and manage the entire stack in the cloud. We are agnostic to a certain number of technologies, which they are not, by definition, and therefore we can optimize the best solution. We can mix and match, and that gives us a significant advantage as well.

And maybe last, but not least, we don't have any legacy to protect, so we can leapfrog to the leading edge and don't have to worry about cannibalizing this or the other part of our legacy software business because in that space we don't have enough.

They don't have a public cloud offering, they don't pretend to have an open marketplace where you can have at the same time consumer and enterprise applications. What we really aim for is that individual within an enterprise, the famous "prosumer." People who want one device on which they can have their private and professional life nicely separated, where they know in confidence that privacy is privacy. And when a company knows that confidentiality and compliance is also guaranteed, we can provide this because we still have a foot — a pretty big foot, actually — on the consumer side of the business. IBM can't provide that.

What more should enterprise IT readers know about the consumer/SMB/ enterprise app store you outlined as part of your new cloud offerings?

It would give CIOs an opportunity to put at the disposal of the users apps that can be easily consumed by employees of the enterprise that have been certified, approved, secured, and were conformed to IT strategy and IT procedures and processes. Some of them can be very large apps, but then you don't really need to put them into an app store. Some of them can be more short-term things. An application to manage your expenses, an application to use your touchpad in order to capture your expenses — you know, scan them with a camera, upload them, and you are done. A whole bunch of things that will make life a lot easier, a lot simpler.

Then of course there are all of the apps that you could use when it comes to analyzing and looking at data, so it becomes a real catalog of capabilities that can be dynamically managed. If something gets corrupted, something gets polluted, you can take it out. You can immediately remove it from all of the devices if you have such a capability — you can bring your things on-stream. It becomes a completely new way of interacting, where I believe CIOs could close the gap in a significant fashion between the old dilemma that you're delivering value for the business users and actually being ready on the IT side.

What do you see as key to competing and winning in the network space against the Junipers and the Ciscos of the world?

The good news is we must be doing something right, because quarter after quarter, we are gaining substantial market share. We have great technology. We cover a lot of space when it comes to networking. Our price-performance ratio must be very optimal because we just "beat the crap" out of the competition.

One of the reasons why we're capable of doing this is not just because our networking gear is so good; it's also because we have this converged infrastructure approach, where people don't just buy networking with storage or service, that you buy the whole solution — which is what they really want. And because it's all optimized internally as well, it has a double-whammy effect.

How would you respond when Cisco CEO John Chambers says the threat from HP is just a lower-priced alternative — it's not a strategic alternative for customers in networking?

With all respect to John, if we can do the same thing at a cheaper price than what he does, why wouldn't that be a strategic alternative?

I think he's talking around things like fabric architectures and the vision of the next-generation data center network.

That's what we're talking about with converged infrastructure, except that we have it. He's still in the PowerPoint version.

How are you weaving 3Com assets into the overall story?

3Com is totally integrated into our networking capabilities. The guys in ESSN [Enterprise Servers, Storage and Networking] under Dave Donatelli's leadership are doing a great job. It is now really selling extremely well as a stand-alone solution, but it fits beautifully into our converged infrastructure as well.

We are quite capable of using our 3Com capabilities when we talk about next-generation data centers that we actually deliver. We have many customers that we are now bringing into the cloud or springing up private clouds in less than 30 days.

With tablets, what will you be offering the enterprise that Apple can't?

Two things: There are a certain number of native things that are built into webOS that made it into a very unique proposition. The best way to describe it is that it's capable of truly multitasking, it's capable of sharing information, and it's able to synergize a lot of the things that are happening in the Web. The reason for that is it's the only operating system that has been designed from the ground up to assume that you're always connected.

We are also capable of securing and managing these devices for an enterprise with our technology. The CIO can switch these things on and off whenever he wants, for any user, and all of the capabilities that are developed with it.

You've talked about the consumerization of IT. Extrapolate out [three] years. What is the impact on IT? How significant a change does it force IT to go through?

It will force corporate IT to have significantly faster innovation cycles. That's going to have a massive impact. It's going to have an impact on all of the applications that are being used. Some of the good old client/server or even older applications will simply not be used anymore by the millennial workers because they won't even want to touch this kind of stuff. Context-aware applications are going to be really important because that's what the consumer is having already today. ■

More from HP's CEO

Read the extended interview online.
tinyurl.com/4z4e4fh

HP's CEO outlines cloud plan

BY ROBERT MCMILLAN, IDG NEWS SERVICE

HP CEO Leo Apotheker unveiled last week a new cloud computing platform that puts the company in competition with Amazon and Google.

In addition to an infrastructure-as-a-service offering, HP will also deliver a marketplace for consumer, small and midsize business, and enterprise applications, Apotheker said.

There will be something for every HP customer in the marketplace, he said. "We'll provide a single open market that integrates consumer, enterprise and developer services," he said, speaking at an event held for press and analysts in San Francisco. The cloud marketplace will include an application store, as well as developer tools and enterprise services and support, he said.

HP was vague on details, but Apotheker billed the services as open and able to support many development languages, and designed to be used by any software maker. "We will only vet the applications for security and interoperability," he said.

HP is launching some of the infrastructure services, "as we speak," Apotheker said. But it will take time to build out a platform-as-a-service offering similar to Microsoft's Windows Azure or Amazon's Elastic Compute Cloud. The platform component of HP cloud will be available by 2012, he said.

HP already has the know-how to build such offerings. It is well-established as a vendor of consumer and data center technology, as well as the middleware software needed to glue different applications together. It has ambitions to be as large as some of the existing well-known cloud service providers.

"If you want to be in the cloud business, it has to be large-scale," Apotheker said. "You have to be able to serve customers everywhere."

But it's unclear whether HP can attract software developers to its new platform and excite consumers and developers in the same way as Google and Amazon.

HP expects to launch its app store next year, Apotheker said. That puts it years behind Apple and Google, who have 350,000 and 250,000 programs in their respective mobile software marketplaces.

Trying to catch up to a head-start like that will be a "challenge" for HP, which is hoping to get developers to write for its own mobile devices, said Mark Fabbi, an analyst with Gartner. "They're a long way behind from that perspective," he said. "You wonder how they can be relevant."

On the other hand, the enterprise cloud computing space is very much up for grabs, Fabbi said. "Enterprises are looking for someone to lead them into this hybrid universe," he said. "That's an area where they have a much more practical opportunity."

Investors have been looking for some reassurance since Apotheker took over from Hurd nearly five months ago. HP's stock hasn't done well since Hurd's departure and financial analysts wonder if Apotheker has what it takes to lead the company forward.

While technology companies such as Google and Apple have seen their profits rise on the basis of their success in the consumer market, HP has been looking for a hit lately. Last month, it announced the HP TouchPad — a tablet that enters a market dominated by Apple's iPad. Based on the webOS software that HP picked up in its 2009 acquisition of Palm, the TouchPad is expected to debut in June.

Apotheker, formerly CEO of SAP, is expected to focus more on the software side of HP's business than Hurd. But the big question is whether Apotheker will be able to make his new company's diverse product groups work together and deliver compelling new products.

"If you look at HP, they've got a lot of good areas, but can they get a multiplier effect because they're HP?" Fabbi said. "If they can't they might as well be separate companies."

It's going to take more than a vague announcement of strategic direction to prove that HP is moving in the right direction, however. Referring to last week's event, Fabbi said, "It was like they were saying, 'Get ready for an announcement, we have some good stuff coming, but we're not really going to tell you the details yet." ■

HP's grand cloud plan

- Developing cloud services from infrastructure-as-a-service offerings to platform services that can be exploited by partners.
- Opening a marketplace for consumer, SMB and enterprise apps, as well as developer tools and enterprise support.
- Building webOS, known as a mobile OS, into its PCs and printers too.

▶ **Rogue,** *from page 1*

We'll call them Mr. North, who is director of network operations for a midsize manufacturing company; Mr. South, an IT administrator in the poultry business; Mr. East, a university systems admin when he was active in this realm; and, Mr. West, a senior systems admin in the medical industry. Here's what they had to say:

How common is this kind of activity within IT departments?

MR. NORTH: It is very common to see this kind of stuff going on. As long as the users don't notice something like slow connection speeds or not being able to get their e-mail, no one really bothers us.

MR. EAST: I hadn't really seen it discussed until this topic came up on *Blue's News*, but it seemed apparent then that most of the old faces I'd seen posting (on that site) for years had also done the same things.

MR. WEST: I would say it is rather commonplace. Obviously at different orders of magnitude depending on how strict management is and the awareness level of people who aren't in on it.

Describe some of the games that you've hosted on company equipment.

MR. SOUTH: I hosted a 24-slot "Counter-Strike: Source" on a company T-1 for about three years. I brought in my own server and put it under my desk and ran it that way. The only company equipment involved was the switch I plugged into and the router that hit the net. I also hosted a 20-person TF2 server for two years during the same period. This was hosted on a decommissioned server that the company wasn't using for anything. ... We mainly played at night. I don't recall any significant activity during the day.

MR. NORTH: Currently I have "test realm" for "World of Warcraft" running that we use to test out gear and specs before we commit to doing so with the actual pay version. I have a Red Hat system that is just used for DNS and mysql server that we are hosting the "WoW" server and vent server on.

MR. WEST: We've had "Team Fortress 2," "Killing Floor," "Counter Strike," "Minecraft" and a few others. We've actually run the servers off a few different boxes. As the company grew/changed we'd need to switch things over to a different box so as not to overload a production box with non-production processes. Obviously it's in our best interest to not cause downtime or other issues so as to not draw attention.

What are the primary motivations for doing this stuff? Saving money?

MR. NORTH: Really it's about two things: The cost savings of hosting our own vent server alone is worth it, but also it's a learning experience for the techs; they have to maintain security at all times on the network as well as load balancing and QoS to allow this to run as smooth as possible.

Mr. West: My motive is to have a free server for myself and my group of friends. We essentially have full control of the box including creating users, running services, compiling code, etc. If we didn't have the free server I highly doubt we'd have one at all. Half of the fun is in flying under the radar.

MR. EAST: A lot of it was "because I could."

How much do you worry about getting caught?

MR. SOUTH: I didn't really worry. I wasn't using bandwidth during peak hours, and I was on great terms with my boss (the CFO).

MR. WEST: It is a mild concern, but by and large such things are allowed (tolerated?) with a wink and a nod. There's also an understanding that the games will not have an adverse effect on business. ... It's hard to get caught when you're the one in charge of the servers and no one else looks at them.

MR. NORTH: I never worry; I mean, that's why we are hired is because no one else can do what we do, and anyone smart enough to find out should come and talk to me about a job!

Did you ever have any close calls where you almost got caught?

MR. NORTH: Yes, it was the result of an office prank where someone attached speakers to a tech's workstation and had them on full. I had the owner of the company in my office and the tech alt tabbed back into a game, which alerted the boss that something was going on. As he got up to go look, I had used VNC to shut down the workstation. I blamed the noise on a PC that was going bad and said that it did that from time to time, which resulted in money to upgrade our workstations. So it was a close call and a blessing all at once.

MR. WEST: We popped the [circuit breaker] in the rack, causing a servicewide outage for about an hour or so. One of the members in the group had acquired a high-power server that would be capable of running dozens of VMs at a time. He offered to let the group use it provided it could be put in the rack with the rest of the servers. We didn't think any harm would come as it would replace the current box we were using.

After plugging the server in and letting it run for a few days, all seemed well. That was until we actually started adding VMs to the machine. The extra VMs increased the load, which increased the power usage, which overloaded the circuit breaker in the rack and brought it down.

MR. EAST: There was never a mention of the game server for the best part of three years, and one day during a staff meeting, I referred to the server by name and my boss said, "Is that the one with the game server on it?" I still have no idea if he was joking, and he certainly didn't care if he wasn't. It was never mentioned again.

Why do you think it's OK to do this?

MR. SOUTH: I never really thought about it in terms of right and wrong. I used company resources that were not being used by the company to build and maintain a community of gamers. I spent lots of time in my office, almost an unhealthy amount. I just saw this as an unspoken benefit of my job.

MR. NORTH: The way I see it, we keep the network running in tip-top shape, we get the job done and no one really ever complains, so why not reward my techs by allowing them to do this? Other people who do well at my work get company cars and different perks, but not us in the IT department, so this is my way of keeping my techs happy.

Also the equipment is never in use (for business purposes) after 5:30 and on weekends, and since we are paying for the bandwidth, we might as well make use of it. ■

ILLUSTRATION: DANIEL VASCONCELLOS

PHOTO BY JOHN HILMAN

Chris Crowell
PRESIDENT AND CEO
ENTERASYS NETWORKS,
A SIEMENS ENTERPRISE
COMMUNICATIONS COMPANY

After 13 years with Cabletron, Crowell leads Enterasys with a forward-thinking technology vision for its wired and wireless solutions and high standards for the customer experience.

FOR MORE INFORMATION:
please visit www.enterasys.com or call Sales at +1 (978) 684-1000 or toll free in North America at (877) 801-7082.

Bring Your Own Device

Managing the Chaos of Mobility and Virtualization.

In the complex world of wired and wireless networks, IT organizations are being pushed to the brink. Now, Enterasys' Crowell offers learned insight into managing the chaos.

If the workforce "goes mobile," how is the enterprise impacted?
There is no "if" about it. The workforce is mobile. Workers need to be able to connect to the corporate infrastructure from any location on any device. And, from a productivity perspective, it's in everyone's best interest to make this happen. Many enterprises worry about how to secure and manage the latest personal devices out there—currently iPad or Droid tablets and smartphones—but Enterasys customers don't struggle with this because our wired and wireless solutions are new-technology ready. We have always been B.Y.O.D. (bring your own device) friendly.

"Workers need to be able to connect to the corporate infrastructure from any location on any device. And, from a productivity perspective, it's in everyone's best interest to make this happen."

How can a perpetually changing mobile environment be managed?
You need management capabilities that provide visibility and control over every asset on the network. You need to know what user, application or device is communicating to what user, application or device—and for what purpose. But IT teams can't financially realize this without an automated, single-pane-of-glass management tool, as it pushes controls throughout the entire infrastructure with the click of a button. This increasing complexity applies to virtual environments now, too.

What's driving that virtualization?
Virtualization allows you to reduce your computing footprint tremendously for cost, resource and space savings. It also brings flexibility and velocity to your ability to deploy and maintain services. In the old days, you had to buy, configure and load machines; but today you can spin them off rapidly. So you need the infrastructure to keep pace. That requires a vendor-agnostic network management strategy that knows what virtual services get created, what can be connected to them, and who is using them.

What must IT organizations demand of network architecture going forward?
Technology never stops changing, which means you have to future-proof investments. You can't possibly do a massive uplift every three to five years and remain profitable. So invest in hardened solutions with a standards-based, vendor-agnostic vision that accommodates any device, application or technology. For example, when VoIP built momentum, many enterprises had to rip and replace their infrastructure, which made the transition cost prohibitive and the deployment cycle insane. But for Enterasys customers, it was a simple firmware upgrade included with maintenance. Technology aside, you need vendors who pay attention after the purchase order is signed. Weigh the post-sales experience during evaluations, like the tenure of the technical assistance team, access to engineering resources, average deployment times, etc. And in-sourced service and support is a must.

Where does Enterasys fit into all of this?
We offer the industry's "first and best," truly integrated wired and wireless solution that helps enterprises deal with the changing IT landscape. Our single-management platform leverages a robust technology patent portfolio to provide built-in automation and visibility and control capabilities from edge to data center. That means you can provision, manage and secure mobile users on mobile devices accessing applications in the cloud or virtual data center(s). With that, we help IT organizations overcome the fluidity and complexity of business to become true enablers.

enterasys® Secure Networks™

NETWORKWORLD
Custom Solutions Group

Mobile payments in U.S. pitting banks vs. telcos

BY ELLEN MESSMER

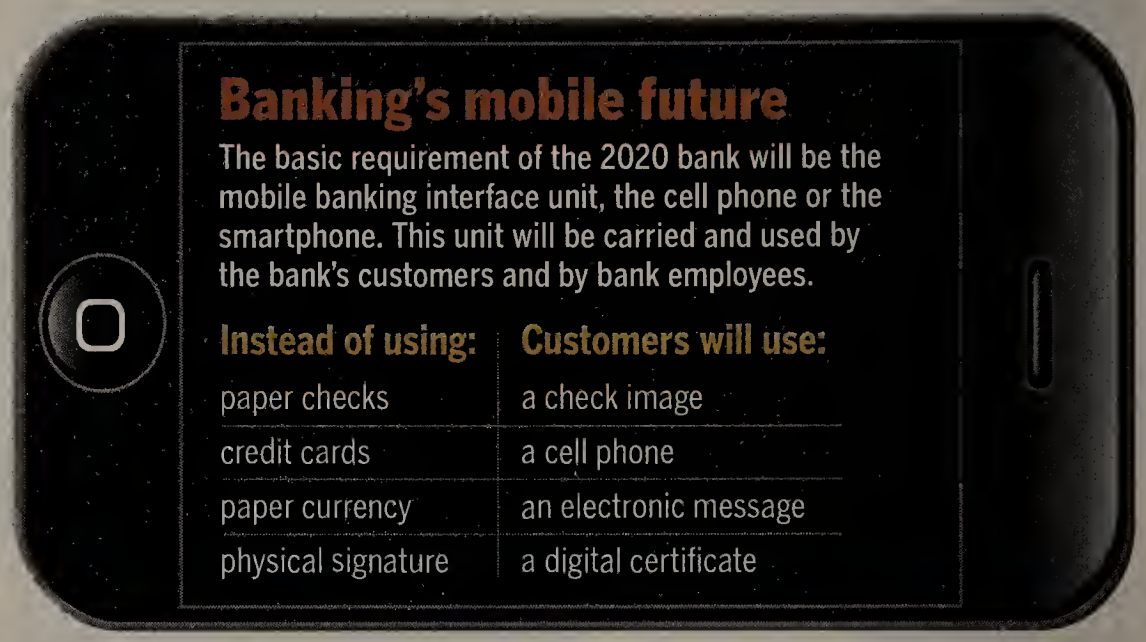

SOURCE: "RETAIL BANK 2020: A ROADMAP TO THE FUTURE," JEROME SVIGALS

THE DAY is nearing when your smartphone will be your wallet, letting you make purchases as stored cash or credit that will be wirelessly accepted by stores or soda machines. Merchants, in turn, will use smartphones like modern point-of-sale devices to process your plastic credit cards.

And smartphones could just "bump" together to transfer money between them.

"Everybody likes the smartphone," says Jerome Svigals, a former IBMer and the author of books predicting the future of banking. "Every major bank in the world has announced a smartphone effort," he says, adding that it appears likely that the wireless "contactless" technology known as near-field communication (NFC) will be a foundation in the new age of mobile payments in the U.S.

But not so fast, say others. While mobile payments in other parts of the world appear to be taking shape through coordination among major wireless carriers selling smartphones, the banks and local retailers, the U.S. sometimes resembles more of a behind-the-scenes brawl.

"The business model has been an active debate," says James Anderson, vice president of mobile-product development at MasterCard, whose constituency is the banks that use its payment-processing services as well as merchants accepting MasterCard. The technology, he says, is not the issue. NFC, which uses the shared 13.56 MHz band, is an ISO standard that MasterCard has backed since about 2005 for mobile-payment use.

"In the U.S., the debate is between the banks and the telcos, and it's an adversarial debate," Anderson says.

The wireless carriers and banks are fighting over transaction revenues and the sense of who "owns" the customer.

The wireless carriers argued "they were bringing tremendous value to the new transaction," whereas the banks argued this is already their payment customer. "And they couldn't find a middle ground," which is slowing innovation, Anderson says.

Part of the fight centers on the Subscriber Identity Module (SIM) card in the smartphone, "a secure element" expected to play a role in managing NFC-based contactless payments, he says. The banks and telcos are at loggerheads, which is why the carriers went off late last year to form their own mobile commerce network called Isis.

Under the Isis banner, T-Mobile USA, AT&T Mobility and Verizon Wireless joined forces late last year to work with Discover Financial Services and Barclays PLC to create a national payment infrastructure for mobile payments based on NFC technology.

NFC is supported today in the Android-based Samsung Nexus S smartphone and is expected to be added to at least some Nokia Symbian and RIM BlackBerry phones. Apple is still leaving everyone guessing about its plans for NFC.

"They want that transaction revenue," says Yankee Group analyst Nick Holland about Isis and its NFC-based mobile-payment network plans. "They are currently working on getting merchants to sign up, but they'll have a hard time."

The obstacle, he says, is likely that Isis presents too closed of a system. "The assumption with Isis is it assumes they own the SIM card and own the transaction," Holland says. But Isis "now seems to be back-peddling" and "talking about more open systems" with a proposal called Open NFC from InsideSecure.

Dave Wentker, head of the mobile product development group at Visa, calls mobile payments "the marriage of card-payment systems and mobile."

Bank of America, Chase, Wells Fargo and US Bank are testing out one type of mobile payment based on MicroSD card functionality, but with "NFC, you need a new phone," Wentker points out.

Nevertheless, work has already produced payment-card terminals that support both NFC and MicroSD cards, he says. Visa points out that 10,000 New York City cabs can accept NFC-based mobile payments and that 200,000 retailers in the U.S. — still a small percentage — have changed their terminals to accept mobile payments.

Last week, a Bloomberg report based on anonymous sources said Google would start testing an NFC-based mobile payment service at stores in New York and San Francisco, paying for the installation of specialized cash registers from VeriFone Systems to accept payments from NFC-based mobile phones.

According to Gartner analyst Avivah Litan, the only mobile-payment system of any magnitude in the U.S. today is the one undertaken by Starbucks in its coffeehouses. It doesn't depend on NFC but a prepaid card for mobile phones based on a barcode system. However, the banking industry is eyeing a day that NFC will be widely used for mobile-payment processing, with Bank of America, JP Morgan, Chase and US Bancorp today all testing terminals supporting NFC.

Decisions on mobile-payment security are expected to be made by the Payment Card Industry Security Standards Council, which sets rules for merchants and processors.

Earlier this year, the council "de-listed" all approved applications for mobile payments that had been included in its PA-DSS certification program — including the VeriFone smartphone-based product for the iPhone. The council says it made this decision to de-list them entirely because it is embarking on a total review of mobile-payment security.

"The rapid development and deployment of these new and innovative mobile payment technologies has brought a level of complexity to the industry never seen before and has introduced a new set of risks and threats that may affect the security of cardholder data," said Bob Russo, general manager of the council. ■

IT Roadmap
CONFERENCE & EXPO

An IDG Enterprise Event

Denver

April 28, 2011

8:15am – 4:30pm

Colorado Convention Center

Register today!

www.itroadmap.net/denad

At **IT Roadmap Conference & Expo**, you'll discover everything you need to know to make informed technology decisions for the year ahead – in just one day. And we're coming to Denver!

Through ROI workshops, technical tutorials, strategy sessions, roundtable discussions, keynote addresses, networking opportunities and an interactive expo floor, you'll hear fresh perspectives and new technology insights from the industry's leading IT analysts, top tech practitioners, and experienced, high-profile end users.

All new morning tracks include:

- **The New Data Center**
- **The Connected Enterprise**
- **The Modern Network**
- **The Public & Private Cloud**
- **The Evolving Threat Landscape**

Who Attends:

- **CIOs & VPs**
- **Directors of IT**
- **IT Managers**
- **Architects & Engineers**

IT ROADMAP 2011

Denver	April 28
Boston	June 7
Dallas	September 13
San Francisco	October
Washington, DC	November

To learn more, visit:

www.itroadmap.net/denad

For more information about sponsorship opportunities and benefits

Contact Andrea D'Amato, Vice President and Publisher of Network World, at adamato@nww.com or 508-766-5455.

PRODUCED BY

CIO COMPUTERWORLD CSO InfoWorld ITWORLD NETWORKWORLD

NETINSIDER | BY SCOTT BRADNER

The congressional view of network neutrality

I'VE BEEN baffled by the inability of Congress to understand the importance of network neutrality.

I'm not much of a fan of regulations for the sake of regulations. There are cases where regulations are warranted, prescription drugs for example, but many other cases where regulations have proven to stop any meaningful progress. Most of the regulations empowering AT&T when it was a monopoly were of the latter type. But I feel that regulations requiring carriers to treat their customers fairly are likely to increase progress rather than limit it.

Some of the people who object to what the FCC is doing claim that the commission does not have the authority under the law to make any rules about Internet network neutrality. This is a legitimate objection.

Others claim that there is not a problem to fix since all the carriers, telephone and cable, have been exemplary Internet citizens and have not done anything bad. This is demonstrably wrong.

But these are not the reactions I'm most concerned with. Too many in Congress, and elsewhere, see that any attempt at ensuring network neutrality will, in the words of Sen. John McCain, R-Ariz., "stifle innovation, in turn slowing our economic turnaround and further depressing an already anemic job market."

This type of reaction only makes sense if someone has absolutely no idea how the Internet works or what it is used for.

The only way such an objection makes sense is if you only look at the carriers and assume that they will be worse off if they cannot get a piece of the action for the business that is done over their networks.

So, the argument must go, let the carriers control everything and they will create jobs and expand the economy.

Let's look at some actual data from the U.S. Census Department. Total U.S. commerce in 2008 (the latest year reported on) was about $22 trillion. Of this about $3.7 trillion was in the form of e-commerce, mostly over the Internet. Most of this (92%) was business-to-business. Doing business over the Internet depends on the Internet working and working fairly.

What about the carriers? The National Cable & Telecommunications Association reports that the total cable company customer revenue for 2008 was about $85 billion and the FCC reports that total U.S. telecommunications industry revenues for 2008 was $297 billion. Thus, total carrier (cable plus telephone) revenue was about $382 billion or about 10% of the value of the business done over the Internet. Commentators that focus on the well-being of the carriers are ignoring the vast majority of the value of the Internet. They want to penalize the 90% to benefit the 10%.

This is an inability to see the value riding over the 'Net, which is the same as having your eyes in your ankles pointing down so they can only see strips of asphalt and miss the cars and trucks riding on the asphalt. But the main problem may be that many of these people can only see "things." They see physical wires and cables but cannot see, so do not recognize, the non-physical traffic using those wires and supporting close to 20% of U.S. commerce.

If you work at a company that uses the Internet to sell to customers or to buy from suppliers you should care about the net neutrality discussion.

Disclaimer: Harvard uses the 'Net a lot but, as far as I know, does not have ankles to house its eyes, or for that matter, eyes to be housed. So the above is my own guess about virtual blindness. ■

New Aruba products blend Wi-Fi, wired access

BY JOHN COX

THE COMPANY that helped pioneer the controller-based enterprise wireless LAN is now selling wired Ethernet switches — with a twist.

Aruba Networks is introducing an array of hardware and software-based services to reshape network access, including a line of wired Ethernet switches that can handle not only the burgeoning array of Wi-Fi clients but also wired clients.

The company last week announced the S3500 series of 24- and 48-port Ethernet switches aimed at the wiring closet. The main difference compared to standard switches: The S3500 automatically seeks out an Aruba Mobility Controller, which handles Aruba's Wi-Fi access points and downloads a set of client policies for configuration, security and management. The switch can apply those policies to Wi-Fi-based laptops, tablets and other mobile clients as well as to desktop PCs or docked laptops.

The line of Aruba controllers is being updated with a new release of the ArubaOS firmware. One change is that the firmware now supports IPv6. Another is Mobile Device Admission Control (MDAC) for Apple iOS. The controller now can identify a device setting up a Wi-Fi connection as an iPhone or iPad, and automatically provision it with certificates, and with security and access policies tailored to these devices. Aruba labels this capability "device fingerprinting."

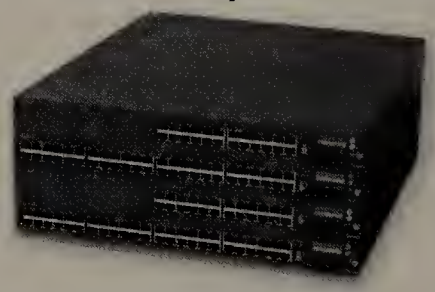

Aruba S3500 Ethernet switch

A related product is Amigopod, which creates an easy-to-use, self-service Web portal that lets employees register their own mobile devices and then get network credentials and access policies tailored to that class of device.

Finally, Aruba is unveiling three access points. First, there are two high-performance 802.11n access points, the AP-134 and -135, both of which have two radios using three data streams, known as 3x3 MIMO, for a maximum data rate of 450Mbps per radio. The AP-134 is outfitted for external directional antennas.

Second is the new Aruba Instant access point family: In a remote site, with a group of these access points deployed, one runs a subset of Aruba's controller software, acting as a controller to the rest of the WLAN. Aruba says the Instant access point can be up and running after a three-minute installation process, and they cluster automatically to receive configurations and updates from their "virtual controller."

Finally, the new AP-175 is Aruba's first outdoor 802.11n product, in a 2x2 MIMO configuration.

Part of Aruba's intent in this announcement is to bring a new level of intelligence about clients to the corporate network, in order to deal with an authorized user who may connect via a wired or wireless connection, and with different devices.

The Aruba Instant AP is expected to be available this month, in two models, priced at $395 and $695. Amigopod also will ship in March. The new S3500 switch, with the ArubaOS 6.1 firmware release, and the other new access points, all ship in April. Aruba will announce product prices then. ■

TOOLS

Gladinet Cloud Desktop, a real cloud product

I wonder how long the rampant marketing hype over cloud-related stuff will last? Now, let's be clear, that's not to say I don't think there's validity in the concept of cloud services; not at all... it's just that many vendors choose to conflate whatever they're selling with the word "cloud" just because it's the *meme du jour*, which does nothing but make the term "cloud" less useful.

Mark Gibbs' Gearhead

A good example of this wanton adoption of "cloudiness" is in the software-as-a-service (SaaS) market. The history of the term "software as a service" goes back to a 2001 article titled "Strategic Backgrounder: Software as a Service" by the Software & Information Industry's (SIIA) eBusiness Division. The phrase was used to describe hosted application services and, unfortunately, it and its ugly acronym were swiftly adopted and became part of IT "industry speak."

But as usual in the IT industry, the dark forces of marketing intervened and over the last few months many SaaS vendors now proclaim themselves to be "cloudy" despite the fact that their products and services are still essentially the same as they were before "cloudiness" got everyone excited.

Now it's true that some SaaS vendors have moved elements of their infrastructure over to be driven on the back end by true cloud services, such as Amazon Web Services, but I'm not convinced that just because a vendor uses a cloud service as part of its offering it too becomes a cloud service (i.e. cloudiness is not a transitive property).

Keith Shaw is off this week. Cool Tools will return in the next issue.

But some vendors do deliver on their cloud promise. For example, I just got my hands on Gladinet Cloud Desktop 3, a product I wrote about just less than a year ago and I'm really impressed with what the company has achieved with this release.

> Gladinet Cloud Desktop ... is simple to configure and manage, it is amazingly useful, it is stable and it is an **excellent value for the money**.

To save you the trouble of reading my previous review, GCD is a Windows utility that maps a drive to a virtual subdirectory under which various cloud services can be configured to appear as subdirectories. Those cloud services include Amazon S3, Synaptic Storage as a Service, EMC Atmos Online, any FTP server, CIFS shares, Google Docs, Mezeo, Rackspace CloudFiles, Windows Live SkyDrive, Windows Azure and WebDav. These cloud services can all be treated like any other Windows accessible storage subsystem, making it very simple to update remote storage resources.

The biggest changes to Gladinet Cloud Desktop are in performance (which appears noticeably faster), the management console (which has been simplified and is somewhat easier to use though it is not as aesthetically "polished" as I had hoped), and the Cloud Sync Folder (which will not only sync a folder between PCs but also supports versioning).

I must also note that Gladinet's marketing people have done one of the things that really annoys me online: They didn't proofread their Web site. Blog postings such as "This ease the need for users that need to both have Cloud Desktop's functionality and also Cloud Backup's snapshot backup functionality for folders and files, SQL Server and etc." make it sound like the text has been badly translated. Given that the company is in Florida, this mangled English is rather odd.

A problem I observed with the previous version —that "if you try to open a document that is on a drive mapped to a cloud services in Word or Excel, all you'll get is an empty document" — appears to have been fixed and I've opened both Word and Excel documents completely painlessly.

There is a perpetually free starter edition available from Gladinet (though they do make it rather hard to find it, tinyurl.com/yb4saks), while the professional version is available at the reasonable price of $59.99.

Gladinet Cloud Desktop is a great product. It provides performance that is as good as the services it accesses, it is simple to configure and manage, it is amazingly useful, it is stable and it is an excellent value for the money.

And, rather refreshingly, it really is a "cloud" product. Gladinet Cloud Desktop 3 gets a rating of 4.5 out of 5. ■

Gibbs watches for clouds in Ventura, Calif. Your observations to gearhead@gibbs.com.

techdebate

EXPERTS FACE OFF *on the* HOTTEST TOPICS

Net neutrality: needed or not?

Innovation begins with an open Internet

M. Chris Riley is policy counsel for Free Press.

NET NEUTRALITY, A FOUNDING PRINCIPLE of the Internet, guarantees that no ISP can dictate where you go and what you do online. Without net neutrality, AT&T, Comcast and Verizon would be free to favor Hulu, but block Netflix. Or prioritize YouTube over Vimeo.

Net neutrality is about protecting the status quo that is the open Internet. On many levels, the term "open Internet" is redundant. If it weren't open, it wouldn't be the Internet. The Internet as we know it isn't about the pipes that connect people; it's about the people that connect over the pipes, the messages delivered over the pipes, and the freedom of all people to use and all messages to travel over those pipes freely.

Clear, enforceable rules are needed to preserve the openness that underlies the Internet because there is no competitive market to protect it. Broadband users typically have only two choices in service providers, and, as their performance expectations rise, many will find only one. The picture for wireless users is little better, thanks to exclusive agreements and early termination fees, among other obstacles to effective competition. Without competition, broadband users can't just choose a less-restrictive service — they are a captive audience, forced to pay what the provider demands for what the provider will allow, or simply go without.

The FCC took a partial step forward when it adopted its Open Internet Order in December, but the rules are riddled with loopholes, fail to include adequate protections for wireless users and fall short of real net neutrality. Yet, despite the rules' weakness, opposition is strong, and the commission is under assault in Congress and in the courts.

Through these challenges, the industry makes clearer every day that it does not intend to preserve the open Internet, but to destroy it. Left to their own devices, the broadband gatekeepers will chisel away at our right to engage in open Internet communications.

Nearly every major Internet business, including Google, Skype, Facebook and Netflix, is the product of the open Internet. None were started by network operators, and all have depended from the start on being able to reach end users over an open connection, without closed gates or toll roads.

► *See* **Riley,** *page 22*

No need for net neutrality regulation

Scott Cleland is chairman of NetCompetition.org.

NET NEUTRALITY REGULATION IS UNNECESSARY, unjustified, unwarranted, unproductive, unwise, unpopular and unlawful.

Net neutrality regulation is unnecessary; it is a solution in search of a problem. Internet users have long enjoyed access to the lawful content of their choice without any government intervention. The FCC's December net neutrality decision is akin to the government regulating all beaches because they found a problem with one or two grains of sand.

Net neutrality regulation is unjustified. The FCC's Open Internet Order included: no market analysis indicating market failure to justify intervention; no assessment of the insufficiency of competition to justify abandoning 15-year-old competition policy; no cost benefit analysis to show the speculative benefits of pre-emptive action would outweigh the real costs of Internet regulation. Ironically, the FCC's net neutrality regulation runs counter to President Obama's January executive order mandating "least burdensome" regulation to promote economic growth and job creation.

Net neutrality regulation is unwarranted. The entire broadband industry fully supports its customers being free to access lawful Internet content of their choice. When the complaint arose about Comcast's use of network management tools that limited BitTorrent, Comcast worked cooperatively with BitTorrent to collaboratively find an acceptable, non-discriminatory and reasonable network management approach. Additionally, the broadband industry created a collaborative engineering working group (BITAG) to resolve network management issues without the need of government involvement, building upon the proven successful model of the Internet Engineering Task Force (IETF).

Net neutrality regulation is unproductive. Historically, communications legislation has been bipartisan. The 1996 Telecom Act, which had the purpose of "promoting competition and reducing regulation," passed Congress near unanimously. That deep bipartisan consensus around promoting competition has been destroyed by some radical net neutrality proponents, who have unproductively polarized large swaths of communications policy by engaging in negative political campaign tactics of demonization and unsubstantiated allegations.

► *See* **Cleland,** *page 22*

Is net neutrality needed?

Yes - 69%

No - 31%

Cast your vote and see comments at
tinyurl.com/4m652jz

▶ **Riley,** *from page 20*

Continued open access to end users remains essential, both for their businesses and for the next Internet start-ups that will rise to challenge them.

The Internet also offers countless benefits for our culture and our democracy. Web sites like YouTube and services like Twitter have opened the door to previously unimaginable possibilities for user participation in the creation and distribution of media — not just its consumption. The open Internet is an antidote to economic and technological restrictions on free speech — it is Gutenberg's printing press on steroids, allowing each and every American with a connection and a computer (or even a phone) to be writer, editor, publisher and reader, all at the same time.

And all of these benefits are at risk if ISPs — instead of users — choose what lawful content, applications and services can be exchanged, offered and utilized. Existing and popular services might become largely inoperable, and new services might never get off the ground, particularly if they compete with services offered by network operators.

The pattern is already on display in the wireless sector. Verizon has offered Android smartphones that come with Google Maps disabled, pushing subscribers toward Verizon's $10 per month navigation service instead. AT&T and other carriers have blocked Skype for years to preserve the revenue from phone service.

The loss of net neutrality will result in significant damage to our economy, our culture and our democracy. By contrast, establishing meaningful rules to protect the right to control your own Internet experience would encourage innovation, participation and competition and enable the United States to regain its status as a global leader in technology and innovation. ■

Free Press is a nonpartisan, nonprofit group working to reform the media.

▶ **Cleland,** *from page 20*

The result has been an unproductive policy climate of controversy that undermines investment, economic growth and job creation.

Net neutrality regulation is unwise. The age old wisdom of the Hippocratic Oath applies here: "First, do no harm." So does the bedrock common-sense notion: "If it ain't broke, don't fix it."

Net neutrality regulation is unpopular. Prior to the midterm election, 302 members of Congress, a majority, wrote to urge the FCC to defer to Congress on net neutrality. In the 2010 midterm election, all 95 candidates that signed a public pledge to support net neutrality regulation lost. Tellingly, net neutrality regulation went 0-95 in the only proxy referendum of the national electorate.

Net neutrality regulation is unlawful. Less than a year ago, the D.C. Court of Appeals ruled in Comcast v. the FCC that the FCC did not have statutory authority to regulate broadband. If the FCC disagreed with that ruling, they should have appealed to the Supreme Court for vindication. Tellingly, the FCC did not. To make matters worse, the FCC's Open Internet Order repeated its previous grievous legal mistake by self-asserting near boundless implicit, or ancillary, legal authority to regulate anything that communications touches. Given that the U.S. Constitution is based on the foundational principle of separation of powers and given that Congress was given the sole constitutional power to legislate, the courts are highly likely to rule the FCC's net neutrality regulations unlawful.

In sum, it is unbelievable that the political debate over net neutrality regulation must continue when net neutrality proponents' arguments are so devoid of merit, justification, evidence, productivity, wisdom, popularity or lawfulness. ■

NetCompetition.org is a pro-competition e-forum supported by broadband interests.

Send Debate Suggestions to jdix@nww.com

The world has changed

Scott, your main argument seems to be, "We haven't had a problem yet, so there will never be a problem." The world doesn't work that way. The world of net access has changed as Internet access became a utility needed to live in the modern world. That transition from luxury to necessity changes the market and is why net neutrality legislation is needed. **— Anon**

Unintended consequences

As with all government regulation, it is the unintended consequences of which we should all be wary. Example: Using corn for fuel which (a) actually pollutes the environment more than gasoline and (b) drives up the cost of food. Example: Forcing banks to make a percentage of home loans to unqualified borrowers, thereby crashing the financial system and wreaking havoc on the housing market. Remember, the elected government representatives who are attempting to implement net neutrality regulations are bought and paid for by special interests! Whether it is big business or public sector unions, the politicos push legislation which favors those who fund their elections. **— Anon**

It always amazes me ...

... when those who complain about technology outstripping the current legal framework bitch and moan when regulations are created to get ahead of the curve. It's not IF these companies will start limiting access to services, but WHEN. Net neutrality regulations would eliminate that uncertainty, assuring the viability of the Internet for years to come. **—Bob R.**

Pro side is weak

Observation: The points made in support of not needing "net neutrality regulation" mostly reference well-known and established facts. The points made in support of regulation, in contrast, are only accepted as established by proponents of the regulation. This spokesman even goes so far as to invent the concept of "founding principles of the Internet" to make his emotional plea. While he claims this type of restriction goes on all the time and that consumer choice is inadequate to guarantee continued open access, the examples he cites are from an industry with open competition (phone service, NOT broadband). In short, the "pro" argument — as presented here — is weak, poorly supported and logically inconsistent. **— Rex**

Cisco sets the bar for mobile security

Cisco integrates always-on client, VPN/firewall and Web security gateway

BY JOEL SNYDER

Cisco has been a leader in remote access VPNs for the past decade, and its latest release, the AnyConnect Secure Mobility Solution, will make both end users and network managers very happy, despite a few rough parts.

The AnyConnect Secure Mobility Solution (part of Cisco's Borderless Networks initiative) consists of three seamlessly integrated products: the AnyConnect Secure Mobility Client 3.0, the ASA Adaptive Security Appliance (firewall/VPN) 8.4 and Cisco IronPort S-series Web security appliance 7.1.

Customers aren't required to buy all three products, but we found that you get better performance and better functionality if you do. Basically, AnyConnect Secure Mobility Solution is all about managed endpoint client software that's always active, protecting enterprise users and enforcing security policy no matter where they are, on a multitude of devices and platforms.

Enterprise network managers will be especially pleased with features such as optimal gateway selection (which automatically picks the best gateway for a user based on network characteristics), endpoint posture assessment and better performance over more diverse types of networks.

It all starts with VPN concentrator

The starting point for any remote access VPN discussion is Cisco's ASA 5500 series Adaptive Security Appliance, a combination VPN and firewall, with optional anti-malware and IPS capabilities.

Although older Cisco VPN clients can connect to non-VPN devices, such as PIX firewalls and IOS routers, connectivity with the new client is more limited. To get the benefit of the AnyConnect client's full feature set, you'll need an ASA appliance. Some IOS routers can also accept AnyConnect clients, but don't support the full feature set.

Your best bet, then, is to use an ASA appliance, which ranges from the ASA 5505 (10 to 25 users) up to the ASA 5585X (5,000 to 10,000 users).

All ASA appliances have SSL VPN features, including reverse proxying (gatewaying Web applications at the application layer) and application tunneling (using encrypted tunnels to expose single applications through the VPN device), although we didn't focus on those features during this test. We spent most of our testing looking at network extension,

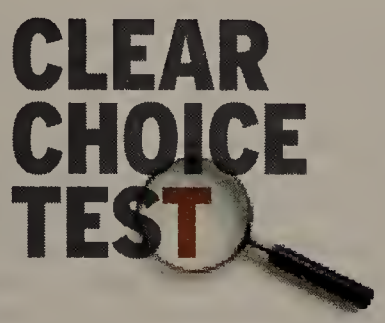

bringing remote devices onto the corporate LAN, and Cisco's approach to securing those remote devices — what is now the traditional remote access use case.

The next piece is Cisco's new AnyConnect Secure Mobility client. The AnyConnect client has the basic feature set expected in a mature product: endpoint security detection and control, simplified deployment and policy downloading directly from the VPN gateway, wide-ranging user authentication options and remote user policy enforcement features.

The AnyConnect client runs on all Windows versions back to XP, Mac OS X 10.5 and 10.6, Intel-based Linux distributions with the 2.6 kernel, Apple iOS 4 (the iPhone and iPad operating system), and Windows Mobile versions 5 and 6.

The AnyConnect VPN client is not required to make a VPN connection to an ASA appliance — you can still use the built-in VPN clients in Windows and Mac OS X, Nokia's Symbian phones, iPhones, iPads and iPods, as well as Cisco's older multiplatform Cisco VPN client, and a host of third-party clients.

However, you give up a lot of performance, functionality and features if you don't use it. For example, the AnyConnect client can use IPSec, SSL/TLS, or DTLS (SSL/TLS run over UDP instead of the normal TCP). We found that shifting from SSL/TLS (TCP) to DTLS (UDP) with the AnyConnect client gave us between 40% and 45% increase in total performance. DTLS and traditional IPSec had similar performance characteristics. Traditional IPSec edged out DTLS by a few percentage points in most of our tests, but the performance difference was difficult to perceive.

Another key feature of the AnyConnect client not found in Cisco's older IPSec clients is endpoint security checking, remediation, and control. Cisco has folded its Cisco Secure Desktop into the AnyConnect client (for a price — there is a license fee), and has merged desktop security management into the VPN concentrator, tremendously simplifying the task of linking desktop and VPN security policies and avoiding the potential for things to drop between the cracks.

Web security is the final piece

The last major piece of Cisco's remote access solution is a new addition: the Cisco IronPort S-series Web Security Appliance. This is a secure Web gateway, with the primary goals of protecting Web-browsing end users from malware and enforcing access controls on where people can browse.

We didn't do a full evaluation of the product, focusing only on its integration with

NETRESULTS

Product	**Cisco Secure Mobility Solution: Adaptive Security Appliance (ASA) 5500-series firewall and VPN concentrator v8.4, IronPort S-series Web security appliance v7.1, AnyConnect Secure Mobility Client v3.0**
Company	Cisco
Price	List price for 250 users: $32,000. That includes ASA 5520 firewall appliance (includes client license) and one year of support ($10,000), plus IronPort S160 Web security appliance with one year of support ($22,000). (Pricing varies depending on configuration and volume discounts are often available.)
Pros	Great end-user experience across multiple platforms; integration of endpoint security and policy enforcement pieces into a single client. Single management pane for most components. Web proxy brings multiple tools, including application controls, into a single device. Automatic integration of ASA and WSA powerful and well done.
Cons	ASA and AnyConnect management complex and hard to learn. Licensing model is too complex.

the ASA and VPN clients. But the IronPort S-series has the expected feature set for a Web security gateway: malware scanning using multiple engines, URL filtering to avoid bad neighborhoods and enforce acceptable use policies, bandwidth management, and the ability to look at content to enforce general security policies, such as blocking PowerPoint attachments.

The IronPort S-series includes "man-in-the-middle" SSL decryption, which lets it scan both encrypted and unencrypted connections, and leverages the IronPort reputation service to do reputation-based lookup of URLs and Web servers. This feature set makes it a fairly complete Web security gateway, not all that different from the other market-leading products.

A cynic might say that Cisco requires network managers to buy a whole separate box — and an expensive one at that — because they don't have built-in Web security in the firewall. That's true, of course, but it's also true that the Web security in the IronPort S-series is more powerful than what you can get with the Web security feature built in to unified threat management firewalls.

Even if you don' t plan to turn on any new features, you'll be happy with the new products because they'll make it easier to do what you're currently doing.

If you already know how to run Cisco's older VPN 3000 GUI, most of the VPN parts have been transplanted into ASDM, Cisco's Java-based ASA appliance management tool Adaptive Security Device Manager.

The ASA appliance can be your source for the VPN client software, and you don't have to build pesky policies that get glued into the AnyConnect client at installation time, so you can have a VPN deployment up and running more quickly than you would using the old client and old hardware.

The AnyConnect client is also more firewall-friendly, falling back to SSL/TLS encryption over the Secure-HTTP (443) port, which means less frustration for end users on the road. And ASDM includes a VPN wizard, to guide you step-by-step and help automatically glue together the bits and pieces that all have to match to make things work.

Legacy licensing

Well, there's actually one problem that will frustrate VPN 3000 users: licensing. The ASA appliance is really the next generation of PIX firewall, with a merging of the best VPN features from both the PIX and the old VPN 3000. One of the features carried over from the PIX is feature-based licensing, which can best be described as "you've got to be kidding."

For remote access alone, there are six types of licenses, with another half-dozen types for the platform itself. For inexplicable reasons, you need a special license to also use mobile devices with your ASA appliance — although only if you use AnyConnect client software, and not if they use the old client, and don't forget the special license for your WSA to make it part of the Secure Mobility Solution.

Fortunately, there's a 48-page manual which explains it all. Our only other advice is to be sure to get your strong encryption license (it's free, fast, and online) before you start, because encryption profiles will only be correctly set up using the wizard if the strong encryption license is already installed.

Putting the pieces together

Cisco Secure Mobility Solution is not just a VPN tool kit; it's about enforcing enterprise security policy when staff members are both in and out of the office. That means you'll need to spend some time thinking about your security policy before you begin configuration.

One of the important things to remember about the AnyConnect client is that it is "always on," meaning that it enforces security policies based on the location of the user, even when there is no tunnel in place. The AnyConnect client periodically connects to the ASA even when the client is not running — you'll see these little 20-packet exchanges to the HTTPS port of the ASA as it verifies that the ASA is alive and well and doesn't have a new policy to hand out.

You can change the security policy on the fly, so you don't have to get it perfect before you start your deployment, but it's a good idea to know where you want to end up before you start. Because the configuration tools within ASDM are so complicated, the only way to avoid getting lost is to zero in on what you want to accomplish. Building policy is only easy to do if you know what you want to enforce.

Cisco could have done a much better job in ASDM of making things consistent and usable. In the VPN part of the GUI alone, there are dozens of options and a confusing and contradictory set of terms. This makes it easy to make mistakes, or build a less secure deployment because you didn't get everything done correctly.

For example, split tunneling can be done with a much higher level of granularity than was available previously, a great security improvement. But digging out the different features and getting them properly configured involves multiple screens and "Advanced" tabs that have to be opened. The result is that it's easier to not use this new feature, and have a less secure deployment.

While much of the VPN feature set can be configured using the command-line interface (CLI), making full use of the feature set requires ASDM. The basic encryption and tunneling tools are all CLI-based, but some parts of the client-side policy configuration rely on hidden files on the internal flash that are best left to ASDM to keep straight.

We built a basic ASA firewall using the CLI, and then we stuck entirely with ASDM. Once we got all of the licensing pieces worked out, our final configuration only took about an hour.

But that was done with the help of a Cisco trainer. The solution has a lot of moving parts, and without hands-on guidance, we could have spent days covering the same territory. If you can afford the time, read through the documentation or take some training.

Happy end users

The good news is that while the Secure Mobility Solution can be complex for network managers, it's a fantastic experience for end users. Think of it as throwing yourself on your sword to help everyone who's actually going to use the remote access VPN. No matter what platform we tested — Mac, Windows and iPhone were in our lab — getting the client installed and operational was

What about IPv6?

We were happy to see good IPv6 support in the AnyConnect client and the ASA appliance. If you give the ASA appliance an IPv6 address on your network and define a pool of IPv6 addresses to hand out, you'll be able to tunnel IPv6 across the IPv4 Internet (although this is only supported in SSL tunnels, not in IKE tunnels). Although you can define IPv6 filters on traffic coming out of VPN tunnels, the AnyConnect VPN client firewall doesn't let you enter IPv6 addresses, so features such as split tunneling aren't fully IPv6-ready. The IronPort S-series has no visible IPv6 support.

— Joel Snyder

CABLESYS
WE KNOW FIBER OPTICS

90's vs. Now

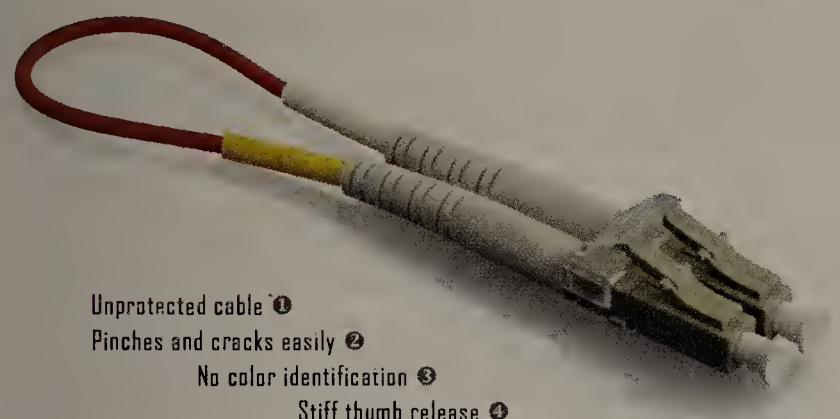

- Unprotected cable ❶
- Pinches and cracks easily ❷
- No color identification ❸
- Stiff thumb release ❹
- Long body ❺
- Flimsy and hard to handle ❻
- Difficult to store ❼

- Protective case ❶
- Ergonomic design ❷
- Color-coded identification ❸
- Flexible thumb release ❹
- Compact body ❺
- Easy to handle ❻
- Easy storage ❼
- Free* ❽

LC Multimode 62.5

LC Singlemode

SC Duplex 50

...and more available.

Still using antiquated loopback adapters? It's time to upgrade.

Fiber optics is our specialty, so when it comes to testing, it is serious business. We designed a compact housing to protect our loopback adapter. Go ahead, put it in your pocket or throw it in your tool box. The fiber core will not bend or crack. We also make it look cool with an ergonomic shape. Best of all, it comes with more than **50%** savings over those big name brands.

*For a limited time, this LC Multimode 50 loopback adapter is available online at cablesys.com/NW for **$7.99** with coupon code, or free with your fiber optic product order over $400.00. Curious? Give your account rep a call to find out more.

*Terms & Conditions: 1. This promotion starts 02/01/11 and expires 05/31/11. 2. Coupon code can only be used online, no call ins. 3. Free offer is not available online, you must call your account rep. 4. Qualifying loopback adapters are GCFAILLM00, GCFAILLL00, GCFAILLS00, GCFAICCM00, GCFAICCL00, GCFAICCS00. 5. This is a direct marketing offer. This offer cannot be combined with any other discounts, promotions or special pricing. 6. Talk to your Cablesys account rep for complete terms and conditions.

800.555.7176 cablesys.com/NW cs@cablesys.com

simple. If end users liked the old Cisco VPN client, they'll love AnyConnect, which has a modern feel and brings benefits beyond just VPN tunnels.

For example, on the Windows platform, AnyConnect client includes Network Access Manager (NAM), a full-fledged 802.1X supplicant for wired and wireless networks. Since AnyConnect client is meant for both the corporate network and roaming, integration of 802.1X features lets a single client handle endpoint security and connectivity.

AnyConnect is your network-access control (NAC) client (with 802.1X and endpoint security checking, remediation, and enforcement) when in the office, and your VPN client (with IPSec and SSL transports, as well as the same endpoint security features) when on the road. Even better, the AnyConnect client can figure out where you are by using a feature called Trusted Network Detection, which looks at domain names and DNS servers being handed out via DHCP. This can help automate the process of choosing whether to use 802.1X and NAC or bring up a VPN tunnel. In our testing using an Enterasys C2 Ethernet switch, Trusted Network Detection and the 802.1X supplicant both worked without any hitches.

It's hard to describe how complete the AnyConnect client experience is without turning this test into a laundry list of features. Cisco has done a good job of covering all the bases, supporting both strict and loose security policies, as well as multiple deployment options and authentication settings. We tried a good assortment of these features and found that in this area the AnyConnect client worked as advertised.

We had mixed success with endpoint security posture checking. Basic host scanning is included as part of the ASA AnyConnect Premium license, while remediation features (such as forcing an anti-malware update or turning on a desktop firewall) require the Advanced Endpoint Assessment license.

Part of the difficulty is that the policy is spread across different parts of ASDM. For example, you look for the presence of a particular antivirus package in one part of ASDM, but you look to make sure you're not executing in a virtual machine in a completely different part of the policy.

The ASDM management tool lets you build a posture checking decision tree using traditional flow-chart symbols. This configuration approach is approximately 10,000% more understandable and scalable than Cisco's old approach based on the ACS RADIUS/TACACS server.

This approach represents Cisco's current thinking on how to do both NAC and VPN posture checking in the same client. Cisco is continuing to avoid the Trusted Computing Group's open standards for posture checking, and has forged ahead with a single-vendor solution, incorporating its own Cisco Secure Desktop and OPSWAT's endpoint posture checking tool kit into a nicely merged solution.

Overall, network managers will have to balance the simplicity of Cisco's strategy, which requires only a single client and no particular cooperation from the endpoint security vendor, with a lock-in to what Cisco and OPSWAT are willing to support.

Our experience with OPSWAT has generally been good, although we have had recurrent difficulties getting consistent results when testing against our lab's standard antivirus package, Sophos. In this test, different configurations of the same antivirus package gave different results in the AnyConnect client. Network managers using the AnyConnect client to do endpoint posture checking will want to experiment with their own configuration and endpoints to avoid false positive and negative results.

Web security goes to the cloud

Cisco's Secure Mobility Solution has three specific strategies for protecting end users from the vast wasteland of the Internet: endpoint security, cloud-based security and enterprise proxy protections.

On the endpoint, the AnyConnect client with its Cisco Secure Desktop feature set doesn't provide much protection itself (beyond a basic personal firewall), but can be used to detect the state of endpoint security and, with an Advanced Endpoint Assessment license, perform some limited controls.

The second strategy, cloud-based security,

Cisco VPN has long history

In 1999, *Network World* tested a dozen VPNs, with a product from Altiga Networks coming in tied for second place. Our main complaint was the lack of split-tunneling capability, a feature that was quickly added.

In 2000, Cisco acquired Compatible Systems and Altiga Networks. The Compatible product, which became the Cisco VPN 500 Series concentrator, was killed off in 2002.

But the VPN 3000 Series from Altiga was an unqualified success. It was easy for end users to work with, supported Windows and Macintosh platforms, and was powerful enough to serve most enterprise remote-access needs. With a range of products from low- to high-end, the VPN 3000 series became the standard for enterprise remote access.

Of the 12 remote-access products we tested in 1999, only two remain on the market: Check Point and Cisco. When we retested VPN client software in 2003, Cisco came out on top of a field of 10 players.

Cisco's domination of the VPN market was so complete that competitors were forced to create a whole new category, SSL VPN, to even think about going up against the VPN 3000 series. The SSL VPN attack has broadened the market for enterprise network managers slightly, with Juniper, F5, and SonicWall as credible alternatives.

But Cisco hit a serious snag in 2005 when it released the ASA 5500 series security appliance, an attempt to merge their successful PIX firewall product line (canceled in 2008) with an even more successful VPN 3000 series (canceled in 2007).

At the same time, Cisco started to merge its many endpoint VPN and security tools. The idea was to fold features from its Host Intrusion Prevention, Desktop Security, 802.1X supplicant, SSL VPN and NAC product lines into a single unified client, the AnyConnect VPN Client (also called AnyConnect Secure Mobility Client).

The hitch for longtime Cisco customers was that Cisco ended support for its PIX and VPN 3000 series products, as the new client doesn't support the older hardware.

The chaos surrounding the ASA 5500 increased when 64-bit Vista hit the streets, an operating system that Cisco wouldn't be able to fully support until 2010. So customers who wanted to simply keep doing basic VPN remote access were forced to replace old — but working — VPN 3000 concentrators with newer ASAs in order to handle end-user operating system upgrades.

While this was necessary from Cisco's point of view to integrate a half-dozen overlapping acquisitions, it remains to be seen whether Cisco customers will forgive them and keep the Cisco VPN solution at the very top of enterprise short lists.

— Joel Snyder

is offered in conjunction with ScanSafe, a recent Cisco acquisition.

Cisco has incorporated the ScanSafe client tool into the AnyConnect client and the ScanSafe policy management tool into ASDM, making the option of deploying cloud-based malware scanning and Web filtering functionality fairly simple. ScanSafe licensing is completely separate from all other Secure Mobility licensing, and ScanSafe is only supported on Windows platforms.

While the integration makes it easy for an enterprise to select cloud-based scanning, we think that most enterprises will see cloud-based scanning versus enterprise proxy protections as an "either/or" choice.

While the AnyConnect Client has a trusted network detection feature, ScanSafe has a similar feature. Rather than combine the two, each runs independently, letting ScanSafe work in a non-AnyConnect environment. Similarly, all of the Web-based security policies established on the IronPort Web proxy are completely independent of the policies set up for ScanSafe; you can't reuse any of the components and you can't easily translate the policy from one to the other.

We chose to focus on the third type of Web security: the Web proxy. Cisco's approach requires a tight linkage between the ASA VPN concentrator and the S-series Web proxy, in order to transfer authentication information to the Web proxy. Making that linkage is very simple — you just put a common port number and shared secret into both devices, click the "test" button, and if everything is correct, you're done.

The ASA sends the username, but not any group membership information, over to the IronPort S-series, so we had to link to our Active Directory to get this information. Once that was settled, we were able to apply user- and group-based Web security policies.

One of the most important parts of the integration between the AnyConnect client, the ASA appliance and the IronPort S-series is the automatic download of proxy information to AnyConnect clients. We tested this with Windows (Internet Explorer), Mac (Safari, Chrome and Firefox) and iPhone systems all running the AnyConnect client, and had seamless experiences browsing through the VPN tunnel, passed to the IronPort S-series proxy and off to the Internet.

The IronPort S-series has a fairly standard set of protections, including URL filtering (for example, blocking gambling sites), malware scanning with two different engines (Webroot and McAfee in our test system), and Web reputation checking, used to block access to known bad Web pages or objects. The IronPort S-series also supports sanctioned man-in-the-middle, a way to "break in" to the SSL conversation by pretending to be the encrypted Web server with a fake public-key infrastructure certificate.

We briefly tested the malware scanning and URL filtering. As with all URL filtering products, we had a very high success rate, but were able to slip through a few URLs in violation of policy. A selection of 10 recent viruses transmitted into our test lab network were all caught by the malware scanner.

SMBs might feel left out in the cold

Cisco's AnyConnect Secure Mobility Solution is a two-box enterprise play that could pose some problems for small and medium-sized businesses.

The Adaptive Security Appliance (ASA) piece of the puzzle delivers firewall and VPN, but not the other security features found in an integrated UTM device. For example, content scanning for malware requires an add-in hardware module and a subscription, as does intrusion prevention.

The problem is that you can only put a single add-in hardware module in any of the appliances, so you have to pick whether you want IPS or anti-malware in your VPN gateway, rather than having the ability to use both as most other UTM firewalls allow.

When the ASA is acting as a firewall, picking one or the other makes sense, because you usually leave anti-malware to end-point software and an anti-spam gateway. When the ASA is acting as a VPN concentrator, however, having both protections is a very attractive defense-in-depth strategy, but the ASA doesn't allow you to do that directly.

In an enterprise environment, Cisco solves this problem by recommending the second box, the full-feature IronPort S-series Web security appliance.

However, the two-box solution could have a side effect of pushing Cisco remote access out of the price range and complexity level appropriate for many small business networks.

— Joel Snyder

We 'like' the Facebook controls

A new feature in the IronPort S-series is application visibility and control. This lets the network manager monitor and block various Web-based applications directly, separately from the URL filtering part of the product. The version we tested is more of a proof-of-concept than a fully baked application visibility tool, with only eight categories, including "Blogging," "Facebook," "IM," "LinkedIn," "Media," "P2P/File Sharing," "Conferencing" and "Social Networking."

These are a bit of a mish-mash of different applications, many of which could be caught by simple URL filtering. However, the idea appears to go beyond the simple block/allow/warn of URL filtering, and get more specific.

For example, Facebook is broken down into 15 subcategories, such as "Facebook Applications: Games" and "Facebook Applications: Education," which would allow you to differentiate different types of Facebook usage, blocking those you don't allow. For example, you can block all Facebook Events, or you could just block posting of events but allow "like" of events. In LinkedIn's controls, you can block the employment section separately from the messaging section, or you can block job searches separately from job postings.

In our testing, the IronPort S-series did exactly what it said it would — identify applications and apply application controls, including bandwidth limits, as a Web proxy. However, it's clear that for this to work, you need a proper configuration.

For example, now that many Facebook users are selecting to encrypt their sessions, you must use the sanctioned man-in-the-middle to decrypt the SSL, or there's no possibility of applying fine-grained application controls. Similarly, if you want to control BitTorrent, you must force the traffic through the proxy.

Overall, the Web security options within Cisco's Secure Mobility Solution give network managers enough choices to provide strong policy enforcement for end users no matter where they are. ■

Snyder is a senior partner at Opus One in Tucson, Ariz. He can be reached at Joel.Snyder@opus1.com.

An old, unsolved problem came back during testing: how to get an end-user browser to actually use the proxy. ***tinyurl.com/4kn88tg***

Force10 delivers fast, dense switch

But extensive tests uncover ASIC-related anomalies, software limitations

BY DAVID NEWMAN

High port density, high throughput and very low latency are bedrock requirements in the data center, and Force10's new S4810 top-of-rack switch delivers on all three counts.

At the same time, Clear Choice testing revealed some limitations in the "merchant silicon" chips increasingly seen in data-center switches. Tests turned up anomalies in cut-through latency, media access control address learning and link aggregation failover handling. The S4810 also turned in mixed results in multicast scalability.

The S4810 is a 1U top-of-rack switch with multiple interface options. It has 48 SFP+ ports for 1G/10G Ethernet (we tested it with 48 10G Ethernet transceivers) and four QSFP+ ports for 40G uplinks. With 10GBase-SR transceivers, the switch drew 202 watts when idle and 219 watts fully loaded.

The switch runs the Force10 Operating System (FTOS), whose command-line interface (CLI) is nearly a clone of Cisco's IOS. Experienced Cisco users will have no trouble configuring and managing this switch.

Although we tested the switch as a Layer-2 data center device, it also supports Layer-3 features, including major IPv4 routing protocols and static routing of IPv6 traffic, via a $2,000 software upgrade.

NETRESULTS

Product	Force10 S-series S4810 top-of-rack switch
Company	Force10 Networks
Price	Switch: $25,000; 10GBase-SR transceivers: $1,800 each.
Pros	High port density, very low and consistent store-and-forward latency.
Cons	Some data center features not yet shipping; odd cut-through latency; MAC address scalability; sluggish multicast learning.

Significantly, the switch does not yet support some key data center protocols, according to a features questionnaire completed by Force10. These include the data center bridging extensions (DCBX), IEEE 802.1Qbb priority-based flow control (PFC), 802.1Qau congestion notification and 802.1Qaz traffic shaping. Force10 says these features are slated for third-quarter 2011 release. (Go online for features questionnaire at tinyurl.com/4eonf2q.)

Unicast performance

We used the same methodology to test the S4810 as in our January 2010 comparison of 10G Ethernet top-of-rack switches. The only difference this time was that we used 48 instead of 24 ports in measuring Layer-2 unicast and multicast performance.

The S4810 puts up solid numbers when it comes to basic unicast traffic handling. It delivers line-rate throughput, regardless of unicast frame size. Better still for delay-sensitive applications, the S4810 offers sub-microsecond average latency when configured in store-and-forward mode. This is one of the first store-and-forward switches we've tested to break the microsecond barrier.

We expected average latency to be lower still in cut-through mode, but that wasn't always the case. For frame sizes of 256 bytes and larger, cut-through latency was significantly higher than the equivalent test in store-and-forward mode. Further, cut-through latency increased with frame length.

Usually cut-through devices are very fast (since they start forwarding a frame before it's fully received, unlike store-and-forward devices which wait until the entire frame is cached before switching it) and they have roughly the same average latency regardless of frame length.

With the S4810, these properties better described the store-and-forward results than cut-through ones (see graphic, next page).

This is partially explained by a characteristic of the Broadcom 56845 application-specific integrated circuit (ASIC) used in the S4810. According to Force10, the chip still acts in store-and-forward mode for frames shorter than 624 bytes, even when set for cut-through operation. This could explain higher cut-through latency for medium-length frames (say, between 256 and 624 bytes), but it's still puzzling why cut-through latency would be higher for longer frames. The testing RFCs require different measurement methods for store-and-forward and cut-through latency, and we checked and rechecked results to verify we'd used the appropriate methods for each. Force10 and other labs also have confirmed this behavior.

Given the latency results, we'd recommend leaving the switch in its default store-and-forward mode. There's a performance advantage for doing so, and users get the extra benefit of error checking that store-and-forward operation provides.

MAC address capacity

Another anomaly appeared in tests of MAC address capacity, which determines how many devices can be attached to a switch. This metric is especially important for virtualization and cloud computing, where virtual machine counts in a single broadcast domain can rise into the tens of thousands.

The S4810's data sheet states its MAC capacity as 128,000; in practice, we found the limit to be slightly lower, averaging 117,145 addresses. The switch ASIC's hashing algorithm accounts for the difference. To save memory and speed lookup times, ASICs store a hash of each MAC address. With a particular set of addresses perfectly matched to a given hashing algorithm, no two hashes will ever overlap or "collide." In practice vendors cannot predict what addresses customers will use, so some collisions are inevitable.

What's more, the actual number of addresses the switch can learn in production is likely to be far lower than 117,000. Typically, address capacity tests are conducted using only three ports. When we configured the Spirent TestCenter traffic generator to offer a set of nearly 100,000 pseudorandom addresses across 48 ports, the switch learned only about 94,000 of these due to hash collisions. Through trial and error, we found that the switch would learn at most around 25,000 addresses without hash collisions when we distributed addresses across 48 ports.

NETWORKWORLD

From the editors of *CIO*, *Computerworld*, *InfoWorld* & *Network World* –

Your Strategic Guide to Operationalizing Cloud

Explore how, when and where to use cloud to enrich your IT portfolio and make your business more agile, with the original editorial content exclusively featured in this Strategic Guide to Operationalizing Cloud.

We know you enjoy and trust the information you get from *CIO*, *Computerworld*, *InfoWorld* and *Network World*. Now look for the digital magazine – a new, targeted information resource that combines these editorial forces!

Delivered straight to your in-box or available online at:
www.networkworld.com/digital_cloud

Debuting March 28, 2011.

Sponsored by

vmware®

To be sure, 25,000 addresses is still a huge number, more than enough for the vast majority of data centers. Then again, some heavy users of virtualization already are pushing above this figure.

Link aggregation fairness

The S4810 allows up to eight ports to be combined into a link aggregation group (LAG) and uses the link aggregation control protocol (LACP) to dynamically add and remove LAG members. We took one LAG member offline, as might occur in the event of a link or transceiver failure, to see how the switch would distribute that port's traffic across remaining members of the LAG.

Traffic distribution was not uniform in this failover test. After we disabled a port, the switch redistributed all of its traffic to the first two ports in the LAG. On a lightly loaded network this wouldn't be a problem, but it could result in oversubscription and frame loss on a heavily loaded LAG. Still, this is an improvement over what we saw on some switches last year, where all traffic was redistributed to just one other LAG member.

As a final test of unicast performance, we checked the S4810 for "forward pressure," a mechanism some switches use to avoid congestion by forwarding frames illegally fast. The S4810 doesn't have that problem. Its clock is set to run at 40 parts per million, faster than Ethernet's theoretical line rate, but that's well within the 100-ppm tolerance allowed in the Ethernet specification.

Multicast performance

We measured the S4810's multicast performance with tests of IGMP group capacity, group join and leave times and throughput and latency. The first two of these stress the switch's control plane via the switch's software and CPU, while throughput stresses the data plane via the ASIC.

Using IGMP snooping, the switch learned 3,000 multicast groups. That's higher than all but one top-of-rack switch tested last year, and a useful figure for trading and videoconferencing applications that require many multicast groups.

The switch's join/leave times were another story. With all receivers subscribed to 989 multicast groups, the S4810 took an average of 21.7 seconds to join each group and 18.3 seconds to leave. That's much higher than most switches in last year's test. The S4810's maximum join and leave times were higher still, at 49.8 and 53.7 seconds respectively. This suggests an overload of the switch's CPU.

More evidence of an overload came in a buffer-overflow message we saw when running this test (and the group capacity test) immediately after a switch reboot. The fact that the switch did not display this message on the second and subsequent test iterations suggests an issue with initial loading of a multicast software module into memory when large group counts are involved. Another issue we saw is that the switch's CLI erroneously reported the same port twice as a member of a given multicast group.

Force10 said it replicated these results in-house, and found much lower join and leave times — of one second or less — when 100 groups were involved instead of nearly 1,000. The vendor also says it's doing more optimization work on this new platform.

The final set of multicast tests examined switch throughput and latency. In these tests, we configured the Spirent TestCenter traffic generator to transmit multicast traffic to one port, and act as multicast subscribers on the 47 remaining ports.

The switch offered line-rate throughput of multicast traffic, with the exception of jumbo frames. With these 9,216-byte frames, the highest zero-loss rate was roughly equivalent to around 98.5 percent of line rate. That's a bit of a surprise in that most data-center switches deliver line-rate throughput in all cases. On the other hand, jumbo frames are more common for unicast than for multicast transport; thus, the multicast jumbo throughput result probably isn't a concern for most users. Average and maximum multicast latencies were roughly comparable to unicast with the switch in store-and-forward mode.

For network managers whose foremost switch requirements are high port density and very low latency, the S4810 is a good fit. The S4810 still has more work to do in the areas of data center features support and multicast processing speeds. These involve software fixes, and Force10 says they're already in the works. The hardware anomalies, such as MAC address learning and link aggregation failover, may take longer to address. ■

Newman is a member of the Network World Lab Alliance and president of Network Test, an independent test lab and engineering services consultancy. He can be reached at dnewman@networktest.com.

Store-and-forward vs. Cut-through

Force10's switch delivered sub-microsecond average latency when configured in store-and-forward mode. Surprisingly, cut-through latency was higher than store-and-forward, so we recommend that cusomers avoid cut-through mode for this switch.

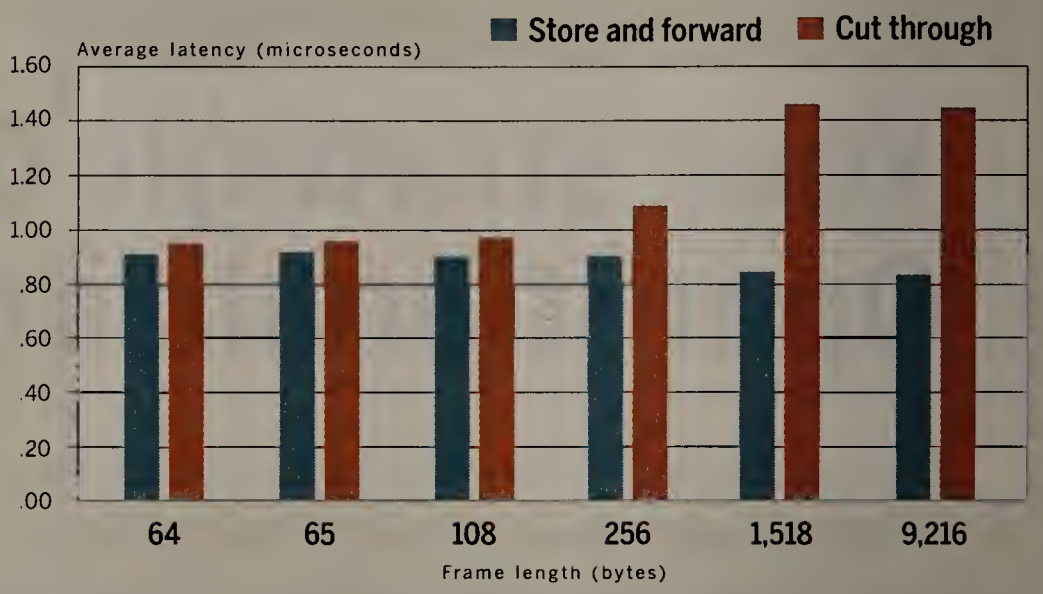

THANKS

Network World gratefully acknowledges the support of vendors that supplied key test bed infrastructure and support. Spirent Communications supplied its Spirent TestCenter traffic generator/analyzer with 10-gigabit Ethernet test ports, and Spirent's Mike Kanada, Timmons Player, Michael Lynge and Chris Chapman provided engineering and logistical support. Fluke Corp. supplied its Fluke 335 TrueRMS clamp meter for power measurement.

NETWORKWORLD

Editorial Index

Advertiser Index

These indexes are provided as a reader service. Although every effort has been made to make them as complete as possible, the publisher does not assume liability for errors or omissions.

*Indicates Regional Demographic

International Data Group
CHAIRMAN OF THE BOARD: Patrick J. McGovern

IDG Communications, Inc.
CEO: Bob Carrigan

Network World is a publication of IDG, the world's largest publisher of computer-related information and the leading global provider of information services on information technology. IDG publishes over 300 computer publications in 85 countries. One hundred million people read one or more IDG publications each month. Network World contributes to the IDG News Service, offering the latest on domestic and international computer news.

Publicize your press coverage in Network World by ordering reprints of your editorial mentions. Reprints make great marketing materials and are available in quantities of 500 and up. To order, contact the YGS Group, (800) 290-5460 ext. 148 or e-mail networkworld@theygsgroup.com.

Network World Events and Executive Forums produces events including IT Roadmap, DEMO and The Security Standard. For complete information on our current event offerings, call us at 800-643-4668 or go to www.networkworld.com/events.

Periodical postage paid at Framingham, Mass., and additional mailing offices. Posted under Canadian International Publication agreement #PM40063731. Network World (ISSN 0887-7661) is published twice monthly by Network World, Inc., 492 Old Connecticut Path, Framingham, MA 01701-9002. **Network World** is distributed free of charge in the U.S. to qualified management or professionals. To apply for a free subscription, go to www.subscribenw.com or write Network World at the address below. No subscriptions accepted without complete identification of subscriber's name, job function, company or organization. Based on the information supplied, the publisher reserves the right to reject non-qualified requests. Subscriptions: 1-877-701-2228. Nonqualified subscribers: $5.00 a copy; U.S.—$129 a year; Canada—$160.50 (including 7% GST, GST#126659952); Central & South America—$150 a year (surface mail); all other countries—$300 a year (airmail service). Digital annual subscription rate of $29.00. Four weeks notice is required for change of address. Allow six weeks for new subscription service to begin. Please include mailing label from front cover of the publication. Network World can be purchased on 35mm microfilm through University Microfilm Int., Periodical Entry Dept., 300 Zebb Road, Ann Arbor, Mich. 48106. **PHOTOCOPYRIGHTS:** Permission to photocopy for internal or personal use or the internal or personal use of specific clients is granted by Network World, Inc. for libraries and other users registered with the Copyright Clearance Center (CCC), provided that the base fee of $3.00 per copy of the article, plus 50 cents per page is paid to Copyright Clearance Center, 27 Congress Street, Salem, Mass. 01970. **POSTMASTER:** Send Change of Address to Network World, P.O. Box 3090, Northbrook, IL 60065. Canadian Postmaster: Please return undeliverable copy to PO Box 1632, Windsor, Ontario N9A7C9.

Network World, Inc.

492 Old Connecticut Path,
Framingham, MA 01701-9002
Phone: (508) 766-5301

To Send E-Mail to NWW Staff
firstname_lastname@nww.com

CEO: Mike Friedenberg
SVP, CHIEF CONTENT OFFICER: John Gallant
SVP, GROUP PUBLISHER: Bob Melk
PUBLISHER: Andrea D'Amato

Online Services
SVP, GM ONLINE OPERATIONS: Gregg Pinsky
NATIONAL DIRECTOR, DIGITAL SALES: Elisa Della Rocco
MANAGER, ONLINE ACCOUNT SERVICES: Danielle Tetreault

Custom Solutions
SVP STRATEGIC PROGRAMS & CUSTOM SOLUTIONS GROUPS: Charles Lee

Events
SVP, EVENTS: Ellen Daly
DIRECTOR OF EVENT OPERATIONS: Deb Begreen

Marketing
VP MARKETING: Sue Yanovitch

Ad Operations
SENIOR PRODUCTION MANAGER: Jami Thompson
ADVERTISING COORDINATOR: Maro Eremyan

Finance
VP FINANCE: Mary Fanning

Human Resources
SVP HUMAN RESOURCES: Patricia Chisholm

Circulation/Subscription
CIRCULATION MANAGER: Diana Turco, (508) 820-8167

IDG List Rental Services
DIRECTOR OF LIST MANAGEMENT: Steve Tozeski
TOLL FREE: (800) IDG-LIST (US only)/
Direct: (508) 370-0822

Sales

Northeast/Midwest/Central
ACCOUNT DIRECTOR, INTEGRATED SALES:
Timothy Keough, (508)766-5475

Southeast/Mid-Atlantic
REGIONAL ACCOUNT DIRECTOR:
Jacqui DiBianca, (610) 971-0808, FAX: (201) 621-5095

Northern California/Northwest
ACCOUNT DIRECTOR, INTEGRATED SALES:
Patti Sellman, (415)267-4522

Silicon Valley/Southwest/Rockies/Utah
ACCOUNT DIRECTOR, INTEGRATED SALES:
Coretta Wright,(415) 267-4515

Marketplace/Emerging Markets–National
ACCOUNT MANAGER, EMERGING MARKETS:
Enku Gubaie,(508) 766-5487

ONLINE

Northeast/Midwest/Central
ACCOUNT DIRECTOR, DIGITAL SALES:
Melissa Rocco, (508)766-5491

East
ACCOUNT DIRECTOR, INTEGRATED SALES:
Jacqui DiBianca, (610)971-0808

Northern California/Northwest/Rockies/Utah:
ACCOUNT DIRECTOR, DIGITAL SALES:
Katie Layng, (415) 267-4518

Northern California/Southwest
ACCOUNT DIRECTOR, DIGITAL SALES:
Katie Albang, (415) 267-4510

Custom Solutions
EASTERN SALES DIRECTOR: Brett Ferry, (508) 988-6941
WESTERN SALES DIRECTOR: Karen Wilde, (415) 267-4512

Event Sales

DEMO
SVP, DEMO: Neal Silverman, (508) 766-5463

IT Roadmap
ACCOUNT DIRECTOR, EVENT SALES, WESTERN REGION:
Brian Klunk, (415) 267-4513
ACCOUNT DIRECTOR, EVENT SALES, EASTERN REGION:
Michael McGoldrick, (508) 766-5459

BACKSPIN | BY MARK GIBBS

Tall tales and 'The Duck Test'

"IF IT looks like a duck, swims like a duck, and quacks like a duck, then it probably is a duck." — *"The Duck Test," by anonymous*

Have you ever received a message from a friend that tells you about something that gets you all riled up?

Many of these messages end up being tall tales. Here's one that is currently circulating: "Education Department officials are threatening school principals with lawsuits if they fail to monitor and curb students' lunchtime chat and evening Facebook time for expressing ideas and words that are deemed by Washington special-interest groups to be harassment of some students."

My path to reading this outrage stoking assertion started when a close friend forwarded a link to an article titled "Big Brother? Feds Order Schools to Monitor Kids Facebook Posts & Lunchtime Chatter" in an online publication called *The Blaze*. My friend added the comment, "I just can't believe this." My friend was right to be suspicious.

It appears that the story was derived (and liberally quoted) from an article on a news-oriented Web site, *The Daily Caller*, titled, "Fed instructs teachers to Facebook creep students," dated March 16, 2011.

This story has, after just over 24 hours circulation, over 65,000 references to it according to a Google search I did for the headline.

What is curious, and rather obvious, if you read the original letter from the Department of Education, is that you won't find any grounds for the claims regarding government pressure to monitor students' Internet use either at school or at home.

The DoE letter is quite obviously intended to frame and discuss the legal and procedural issues surrounding the problem of bullying and what schools are required to do to address the problem without any specific focus on Internet anything.

Nowhere does the letter say, as *The Daily Caller* article by Neil Munro contends, "Under the new interpretation, principals and their schools are legally liable if they fail to curb 'harassment' of students, even if it takes place outside the school, on Facebook or in private conversation among a few youths." The only "new interpretation" I can find is that provided by Munro and *The Daily Caller.*

The Daily Caller article conflated the DoE letter with a whole barrage of unsubstantiated claims, such as: "There has only been muted opposition to this far-reaching policy."

Really? What muted opposition? How is the "policy" (which isn't actually a policy) "far-reaching" when there's nothing extraordinary in the letter's content! I could go on slicing and dicing, but there's no reason to; the whole tale is baseless and shameful on the part of *The Daily Caller, The Blaze*, and every other online publication that accepted what was unfounded, unprofessional opinion and recycled it as fact.

But we, as online readers, need to be far more critical and more demanding. We need to look for veracity. We need to demand support for assertions of any kind, but especially those that appear to fail the smell test. And we need to make the smell test more rigorous.

When you find yourself thinking, "This can't be real?!" listen to what your common sense is telling you: If it quacks, it's probably not an eagle. ■

Gibbs can hear the sound of ducks in Ventura, Calif. Your common sense to backspin@gibbs.com.

NETBUZZ | BY PAUL MCNAMARA

If you bought 100 shares 25 years ago ...

IT'S MARCH 13, 1986: Microsoft, founded more than a decade earlier and already a powerhouse in the world of personal computer software, executes an initial public stock offering that will raise $61 million for the company and leave 30-year-old co-founder Bill Gates unfathomably wealthy.

If you had the good fortune to have bought 100 shares at the $21 offering price that day and sat on the investment for 25 years, it would have mushroomed into 28,800 shares over the course of nine stock splits and be worth about three quarters of a million dollars today (excluding dividends).

That's the good news. Here's the disheartening caveat: Had you instead sold your stash on Dec. 1, 1999, when Microsoft's stock price reached its peak, you would have reaped $1.4 million.

You have to believe someone did ... and tells that story every day.

Speaking of good fortune, *Fortune* magazine was granted inside access to Gates, his executive team, and their Wall Street partners in the months leading up to the IPO. That arrangement resulted in a terrific fly-on-the-wall story published four months later. Here are a few highlights gleaned from that story and other online resources:

Gates was not anxious to go public, but Microsoft was bumping up against federal regulations governing the number of private stockholders a company can have before being required to register with the SEC.

A quote from Gates: "The whole process looked like a pain, and an ongoing pain once you're public. People get confused because the stock price doesn't reflect your financial performance. And to have a stock trader call up the CEO and ask him questions is uneconomic — the ball bearings shouldn't be asking the driver about the grease."

Crafting the prospectus was reportedly a labor of dental surgery, as the driving goal became guarding against future litigation that might be fueled by even the slightest hint that Microsoft was hyping its future prospects. Look which current CEO pops up as the voice of doom and gloom in a description of one meeting with the Wall Streeters:

"For 10 hours Gates, (Microsoft president and COO Jon) Shirley, and other managers exhaustively described their parts of the business and fielded questions. Surprisingly, the Microsoft crew tended to be more conservative and pessimistic than the interrogators. Steven A. Ballmer, 30, a vice president sometimes described as Gates's alter ego, came up with so many scenarios for Microsoft's demise that one banker cracked: 'I'd hate to hear you on a bad day.'"

And here's how the *Fortune* story described the opening bell:

"At 9:35 Microsoft's stock traded publicly on the over-the-counter market for the first time at $25.75. Within minutes Goldman Sachs and Alex Brown exercised their option to take an extra 300,000 shares between them. (Microsoft CFO Frank) Gaudette could hardly believe the tumult. Calling Shirley from the floor, he shouted into the phone, "It's wild! I've never seen anything like it — every last person here is trading Microsoft and nothing else.'"

Gates earned a mere $1.6 million for shares he sold that day, but his remaining 45% stake in the company was worth $350 million, instantly making him one of the nation's 100 wealthiest individuals.

He splurged by paying off his $150,000 home mortgage. ■

Comments and stock tips to buzz@nww.com.

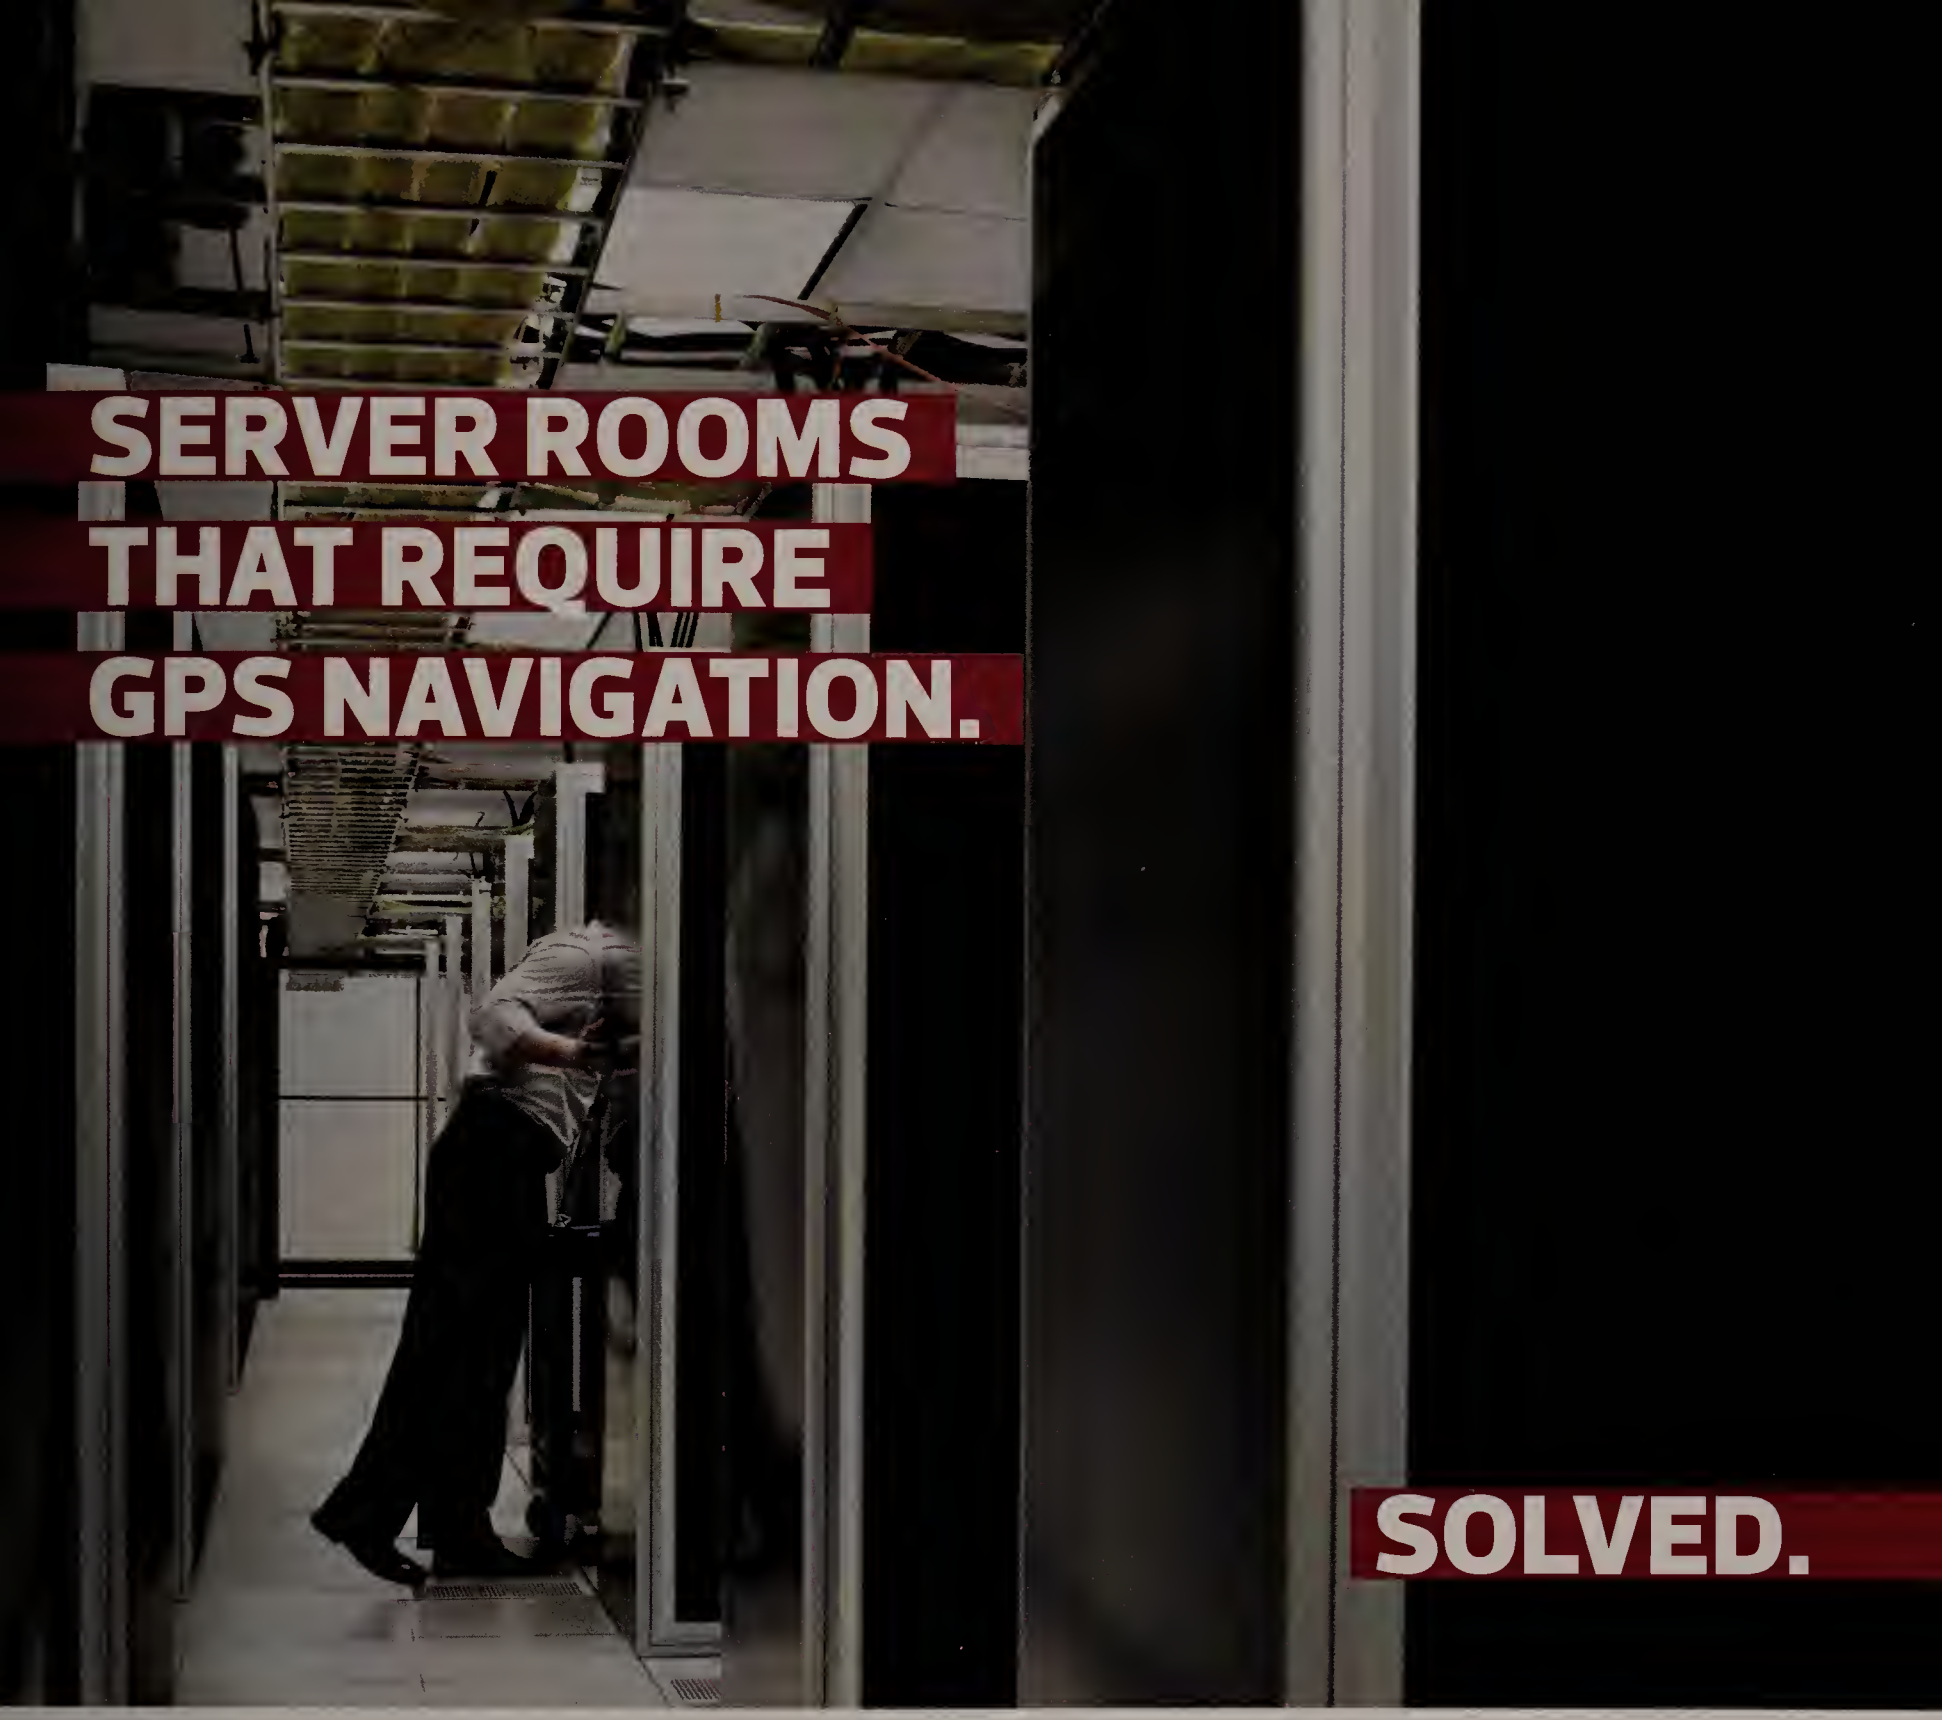

Aaaah, the sprawling server room. It's familiar territory. And we have the innovative technology to get you out. The HP Proliant portfolio is virtualization-ready and will lead the way to data center efficiency. Our HP-trained solution architects will help get you the most out of your investment. And before you know it, you'll arrive at your virtualization destination.

Get double the guidance at CDW.com/hp

hp + CDW PEOPLE WHO GET IT

Microsoft

Now we're talking private cloud, not just virtualization.

Windows Server is changing the conversation.

The virtualized server is a big deal.

It helps businesses, big and small, make IT more efficient.

But what comes next?

Enter the private cloud — a way to manage your infrastructure as a pool of computing resources to deliver your applications and best serve the ever-changing needs of your business.

Windows Server Hyper-V and System Center put you in control with complete end-to-end service management, as well as the ability to tap into the power of the public cloud.

And that's really the whole point of having a private cloud in the first place — control.

With the ability to run and manage multiple hypervisors, development languages, and infrastructure platforms, Microsoft transcends those barriers to put the control in your hands.

Because the technology and vendors you use are there to serve your business needs, not the other way around.

IT is no longer just about hardware. Or software. Or maintenance. It's about finding new efficiencies and new ways of doing things that help your company's bottom line.

So the less company brainpower you devote to fixing old things, the more you can dedicate to coming up with new things.

More computing power. And more available brainpower.

That's Cloud Power.

Microsoft.com/cloud/privatecloud

Cloud Power

Windows Server Hyper-V